HP IO Accelerator Version 3.2.3 Linux User Guide

Abstract

This document describes software requirements for all relevant HP IO Accelerators using Linux operating systems. This document is intended for system administrators who plan to install and use HP IO Accelerators with a Linux operating system. It is helpful to have previous experience with HP IO Accelerators and a Linux operating system. This user guide is intended for IO Accelerator software release 3.2.3 or later.

Part Number: 647094-003
February 2013
Edition: 3

© Copyright 2011, 2013 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

AMD is a trademark of Advanced Micro Devices, Inc.

Windows® is a U.S. registered trademark of Microsoft Corporation.

Contents

About this guide 6
Contents summary 6
Introduction 7
Overview 7
Product naming 7
Performance attributes 8
Required operating environment 9
Supported firmware revisions 9
Supported hardware 9
Software installation 12
Installation overview 12
Installing RPM packages on SUSE, RHEL, and OEL 12
Building the IO Accelerator driver from source 14
Building an RPM installation package 14
Upgrading device firmware from VSL 1.x.x or 2.x.x to 3.x.x 15
Upgrading procedure 16
Loading the IO Accelerator driver 18
Controlling IO Accelerator driver loading 18
Using the init script 19
Mounting filesystems 20
Handling IO Accelerator driver unloads 20
Setting the IO Accelerator driver options 20
Using module parameters 20
One-time configuration 21
Persistent configuration 21
Upgrading the firmware 21
Enabling PCIe power 22
Using the device as swap 22
Using the Logical Volume Manager 22
Configuring RAID 23
RAID 0 23
RAID 1 25
RAID 10 25
Understanding Discard (TRIM) support 26
Discard TRIM on Linux 26
Setting up SNMP for Linux 27
SNMP details for Linux 27
Files and directories 27
SNMP master agent 27
Launching the SNMP master agent 27
Configuring the SNMP master agent 28
SNMP agentX subagent 28
Installing the SNMP subagent 28
Running and configuring the SNMP subagent 29

Manually running the SNMP subagent 29
Subagent log file 29
Using the SNMP sample config files 30
Enabling SNMP test mode 30
Troubleshooting SNMP 33
Supported SNMP MIB fields 33
Maintenance 35
Maintenance tools 35
Device LED indicators 35
HP IO Accelerator Management Tool 35
Command-line utilities 35
Enabling PCIe power override 36
Enabling the override parameter 37
Common maintenance tasks 37
Unloading the IO Accelerator driver 38
Uninstalling the IO Accelerator driver RPM package 38
Uninstalling the IO Accelerator Utilities, IO Accelerator Management Tool 2.x, and other support packages 38
Disabling auto attach 38
Unmanaged shutdown issues 39
Disabling the driver 39
Utilities 40
Utilities reference 40
fio-attach 40
fio-beacon 41
fio-bugreport 41
fio-detach 42
fio-format 43
fio-pci-check 44
fio-snmp-agentx 45
fio-status 45
fio-sure-erase 47
fio-update-iodrive 49
Monitoring IO Accelerator health 52
NAND flash and component failure 52
Health metrics 52
Health monitoring techniques 53
About flashback protection technology 54
Software RAID and health monitoring 54
Performance and tuning 55
Introduction to performance and tuning 55
Disabling DVFS 55
Limiting APCI C-states 55
Setting NUMA affinity 56
Setting the interrupt handler affinity 56
NUMA configuration 57
Introduction to NUMA architecture 57
NUMA node override parameter 57
Advanced configuration example 57

Resources 59
Subscription service 59
For more information 59
Regulatory information 60
Safety and regulatory compliance 60
Turkey RoHS material content declaration 60
Ukraine RoHS material content declaration 60
Warranty information 60
Support and other resources 61
Before you contact HP 61
HP contact information 61
Customer Self Repair 61
Acronyms and abbreviations 69
Documentation feedback 71
Index 72

About this guide

Contents summary

- Instructions on downloading and installing the approved driver and utilities
- Instructions on maintaining the IO Accelerator
- Description of the following IO Accelerator models:
  - HP IO Accelerator for BladeSystem c-Class
  - HP PCIe IO Accelerator
  - HP PCIe IO Accelerator Duo

△ **CAUTION:** Before upgrading to 3.x.x software and firmware, back up all data on the IO Accelerator. The 3.2.3 software and firmware reformat the drive, which causes data to be lost if not backed up. The 3.2.3 software is not backward compatible with 1.2.x or 2.x software.

Introduction

Overview

Designed around ioMemory, a revolutionary storage architecture, HP IO Accelerator is an advanced NAND flash storage device. With performance comparable to DRAM and storage capacity on par with hard disks, the IO Accelerator increases performance so that every server can contain internal storage that exceeds the I/O performance of an enterprise SAN.

HP IO Accelerator is the first data accelerator designed specifically to improve the bandwidth for I/O-bound applications.

In addition to the hardware driver, the IO Accelerator also includes a VSL. This hybrid of the RAM virtualization subsystem and the disk I/O subsystem combines the best features of both systems. VSL functions as a disk to interface well with block-based applications and software, while also running like RAM underneath to maximize performance. This feature produces the following benefits:

- **Performance:** The VSL offers direct and parallel access to multiple CPU cores, enabling near linear performance scaling, consistent performance across different read/write workloads, and low latency with minimal interruptions and context switching.
- **Extensibility:** The VSL enables flash-optimized software development, making each IO Accelerator module a flexible building block for creating a flash-optimized data center.

Product naming

HP IO Accelerator Generation 1 devices include:

- AJ876A: HP 80GB IO Accelerator for BladeSystem c-Class
- AJ877A: HP 160GB IO Accelerator for BladeSystem c-Class
- AJ878A: HP 320GB IO Accelerator for BladeSystem c-Class
- AJ878B: HP 320 GB IO MLC Accelerator for BladeSystem c-Class
- BK836A: HP 640GB IO MLC Accelerator for BladeSystem c-Class

IMPORTANT: Generation 1 IO accelerators for BladeSystem c-Class are only compatible with G7 and earlier server blades.

- 600278-B21: HP 160GB Single Level Cell PCIe ioDrive for ProLiant Servers
- 600279-B21: HP 320GB Multi Level Cell PCIe ioDrive for ProLiant Servers
- 600281-B21: HP 320GB Single Level Cell PCIe ioDrive Duo for ProLiant Servers
- 600282-B21: HP 640GB Multi Level Cell PCIe ioDrive Duo for ProLiant Servers
- 641027-B21: HP 1.28TB Multi Level Cell PCIe ioDrive Duo for ProLiant Servers

HP IO Accelerator Generation 2 devices include:

- QK761A: HP 365GB IO MLC Accelerator for BladeSystem c-Class

- QK762A: HP 785GB IO MLC Accelerator for BladeSystem c-Class
- QK763A: HP 1.2 TB IO MLC Accelerator for BladeSystem c-Class

IMPORTANT: Generation 2 IO accelerators for BladeSystem c-Class are only compatible with Gen8 and later server blades.

- 673642-B21: HP 365GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673644-B21: HP 785GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673646-B21: HP 1205GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673648-B21: HP 2410GB Multi Level Cell G2 PCIe ioDrive2 Duo for ProLiant Servers
- 721458-B21: HP 3.0TB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers

Performance attributes

IO Accelerator capacity Models AJ878B and BK836A	**320GB**	**640GB**
NAND type	MLC (Multi Level Cell)	MLC (Multi Level Cell)
Read Bandwidth (64kB)	735 MB/s	750 MB/s
Write Bandwidth (64kB)	510 MB/s	550 MB/s
Read IOPS (512 Byte)	100,000	93,000
Write IOPS (512 Byte)	141,000	145,000
Mixed IOPS* (75/25 r/w)	67,000	74,000
Access Latency (512 Byte)	30 µs	30 µs
Bus Interface	PCI-Express x4	PCI-Express Gen1 x4

IO Accelerator capacity Models QK762A and QK763A	**785GB**	**1.2TB**
NAND type	MLC (Multi Level Cell)	MLC (Multi Level Cell)
Read Bandwidth (1MB)	1.5 GB/s	1.5 GB/s
Write Bandwidth (1MB)	1.1 GB/s	1.3 GB/s
Read IOPS (Seq. 512 Byte)	443,000	443,000
Write IOPS (Seq. 512 Byte)	530,000	530,000
Read IOPS (Rand. 512 Byte)	141,000	143,000
Write IOPS (Rand. 512 Byte)	475,000	475,000
Read Access Latency	68µs	68µs
Write Access Latency	15µs	15µs
Bus Interface	PCI-Express Gen2 x4	PCI-Express Gen2 x4

*Performance achieved using multiprocessor enterprise server

- Enterprise data integrity
- Field upgradeability
- Green footprint, 7.5W nominal per device

NOTE: MSI was disabled to obtain these statistics.

Required operating environment

The HP IO Accelerator with software 3.2.3 is supported for use in the following operating environments:

- Red Hat Enterprise Linux 5 (AMD64/EM64T)
- Red Hat Enterprise Linux 6 (AMD64/EM64T)
- SUSE LINUX Enterprise Server 10 (AMD64/EM64T)
- SUSE LINUX Enterprise Server 11 (AMD64/EM64T)

CAUTION: Version 3.1.0 or greater of the driver software is not backward-compatible with any previous driver version. When you install version 3.2.3, you cannot revert to any previous version.

IMPORTANT: All operating systems must be 64-bit architecture.

NOTE: IO Accelerators cannot be used as hibernation devices.

Supported firmware revisions

After February 19, 2013, all IO Accelerators ship with firmware version 7.1.13.109322 or higher. This firmware version only works with VSL 3.2.2 or higher. If you are installing a recently purchased or a replacement IO Accelerator into a system that already has IO Accelerators installed, then you must upgrade the firmware on the previously installed devices to 7.1.13.109322 or higher. You must also upgrade the VSL to 3.2.2 or higher. Upgrading the firmware and VSL on cards that were running firmware versions 6.x.x or higher and VSL 3.x.x or higher is NOT data destructive. However, HP recommends that you back up any data on the device before performing the upgrade. The latest supported version of the firmware and VSL can be found on the HP website (http://www.hp.com).

Release	Firmware revision
1.2.4	17350
1.2.7	36867 or 42014
1.2.8.4	43246
2.2.x	43674
2.2.3	101583
2.3.1	101971 4 or 101971_6
3.1.1	107004 or greater
3.2.3	109322

Supported hardware

HP IO Accelerator for BladeSystem c-Class

BladeSystem c-Class IO Accelerators have two distinct designs for the respective server product lines. The G1 through G7 IO Accelerator adapter is provided in a c-Class Type 1 Mezzanine card form factor. It can be installed in both Type 1 and Type 2 mezzanine slots within the c-Class blade G1 through G7 servers, enabling a total of two cards in a half-height server blade, and three cards in a full-height server blade and up to 6 in a double-high, double-wide server (BL680c).

The Gen8 adapter is provided in a c-Class Type B Mezzanine card form factor. It can only be installed in Type B mezzanine slots within the Gen 8 or later servers, enabling one IO Accelerator in a half-height Gen8 server.

The Type I mezz card and the Type B mezz card are distinguished by the mezzanine connector. The type B card is slightly larger than a Type I mezz card.

The amount of free RAM required by the driver depends on the size of the blocks used when writing to the drive. The smaller the blocks, the more RAM is required. The table below lists the guidelines for each 80GB of storage. For the latest information, see the *QuickSpecs* sheet to the *HP IO Accelerator for HP BladeSystem c-Class* at HP Customer Support (http://www.hp.com/support).

The Remote Power Cut Module for the c-Class blade mezzanine card provides a higher level of protection in the event of a catastrophic power loss (for example, a user accidentally pulls the wrong server blade out of the slot). The Remote Power Cut Module ensures in-flight writes are completed to NAND flash in these catastrophic scenarios. Without the Remote Power Cut Module, write performance is slower. Writes are not acknowledged until the data is written to the NAND module, thereby slowing performance. When the Remote Power Cut Module is installed, writes are acknowledged by the IO Accelerator controller to the driver. The IO Accelerator controller then completes the write to the NAND module.

The IO Accelerators (QK761A, QK762A, and QK763A) for Gen 8 BladeSystem c-Class have the power cut functionality embedded on the card. They offer the protection without requiring the remote power cut module.

> **NOTE:** The Remote Power Cut Module is used only in the AJ878B and BK836A models. Without the Remote Power Cut Module, write performance is slower.

HP PCIe IO Accelerator minimum requirements

- An open PCI-Express slot—The accelerator requires a minimum of one half-length, half-height slot with a x4 physical connector. All four lanes must be connected for full performance. HP PCIe IO Accelerator Duo requires a minimum of a full-height, half-length slot with an x8 physical connection. If your system is using PCI 1.1, all x8 signaling lanes must be connected for full performance. If your system is using PCI 2.0, for full performance you only have to connect x4 signaling lanes.

  > **NOTE:** For PCIe IO Accelerators, using PCIe slots greater than x4 does not improve performance.

  > **NOTE:** The power cut feature is built into PCIe IO Accelerators; therefore, no Remote Power Cut Module is necessary.

- 300 LFM of airflow at no greater than 50°C. To protect against thermal damage, the IO Accelerator also monitors the junction temperature of its controller. The temperature represents the internal temperature of the controller, and it is reported in `fio-status` report. The IO Accelerator begins throttling write performance when the junction temperature reaches 78°C. If the junction temperature continues to rise, the IO Accelerator shuts down when the temperature reaches 85°C.

  > **NOTE:** If you experience write performance throttling due to high temperatures, see your computer documentation for details on increasing airflow, including fan speeds.

- Sufficient RAM to operate—The amount of RAM that the driver requires to manage the NAND flash varies according to the block size you select when formatting the device (filesystem format, not low-level format). For a virtual machine using an IO Accelerator directly (using PCI pass-through), consult the user guide for the installed operating system. The following table lists the amount of RAM required per 100GB of storage space, using various block sizes. The amount of RAM used in driver version 3.0 is significantly less than the amount used in version 1.2.x.

Average block size (bytes)	RAM usage for each 80 GB IO Accelerator (Megabytes)	RAM usage for each 100 GB IO Accelerator (Megabytes)	Minimum system RAM requirement for 320 GB MezzIO Accelerator*
8,192	250	280	1 GB
4,096	400	530	1.6 GB
2,048	750	1,030	3 GB
1,024	1,450	2,000	5.8 GB
512	2,850	3,970	11.4 GB

Average block size (bytes)	Minimum system RAM requirement for 640 GB MezzIO Accelerator*	Minimum system RAM requirement for 785 GB MezzIO Accelerator*	Minimum system RAM requirement for 1.2 TB MezzIO Accelerator*
8,192	2 GB	2.2 GB	3.4 GB
4,096	3.2 GB	4.2 GB	6.4 GB
2,048	6 GB	8.1 GB	12.4 GB
1,024	11.6 GB	15.7 GB	24 GB
512	22.8 GB	31.2 GB	47.6 GB

* For IO Accelerator use only. Additional RAM is needed for system operating system and applications.

HP PCIe IO Accelerator Duo requirements

In addition to the IO Accelerator cooling and RAM requirements listed in the previous table, the IO Accelerator Duo requires at least:

- A PCIe Gen1 x8 slot or a PCIe Gen2 x4 slot
- A minimum of a full-height, half-length slot with a x8 physical connection. For systems with PCI 1.1, all eight signaling lanes must be active for full IO Accelerator Duo performance. For systems with PCIe 2.0, only four signaling lanes must be active for full performance.

  NOTE: With driver version 3.1.0 and later, the driver detects in the BIOS if the PCIe slot supports a 75W power draw. If the slot supports up to 75W, the IO Accelerator device draws up to that amount of power. However, if an external power cable is used, power is only supplied by that cable.
  To verify whether a slot is supplying 75W, view the system logs or use the `fio-pci-check` utility.

Software installation

Installation overview

For the system requirements, including supported operating systems, consult the *HP IO Accelerator Release Notes* for this release.

Before installing the IO Accelerator driver, make sure you have properly installed the IO Accelerator devices. For more information, see the *HP IO Accelerator Hardware Installation Guide*.

CAUTION: This version of the IO Accelerator driver is required for newer IO Accelerator devices, including IO Accelerator Gen2 devices, to function properly. For more information, consult the release notes for this release.

IMPORTANT: All commands require administrator privileges. To run the commands, log in as `root` or use `sudo`.

1. If necessary, uninstall the previous version of the driver and utilities. For more information, see "Common Maintenance Tasks (on page 37)."
2. Install the latest version of the driver. You can install the driver as a pre-compiled binary package or as a source-to-build package.

   NOTE: To determine whether pre-compiled binary packages are available for your kernel version or to build the driver package from source, follow the instructions under "Installing RPM Packages ("Installing RPM packages on SUSE, RHEL, and OEL" on page 12)."
3. Install the utilities and management software (included in the driver installation instructions).
4. Load the driver ("Loading the IO Accelerator driver" on page 18).
5. Set the options ("Setting the IO Accelerator driver options" on page 20).
6. If necessary, upgrade the firmware ("Upgrading the firmware" on page 21) to the latest version.

Installing RPM packages on SUSE, RHEL, and OEL

1. Install a version of the IO Accelerator software that is built for your kernel. To determine what kernel version is running on the system, at a shell prompt, use the following command:

   ```
   $ uname -r
   ```
2. Compare the kernel version with the binary versions of the VSL software available from the HP website (http://www.hp.com/support):
   a. Search for a binary version of the software that corresponds to the system kernel version and download it. For example:

      ```
      iomemory-vsl-<kernel-version>-server_<VSL-version>.x86_64.rpm
      ```
   b. If a binary version of the software corresponding to your kernel is unavailable, then download the source `rpm` package. For example:

      ```
      iomemory-vsl_<VSL-version>.src.rpm
      ```

IMPORTANT: Exact package names vary, depending on the software and kernel version chosen.

3. Go to the HP website (http://www.hp.com/support), and then download all of the support `rpm` packages. These packages provide utilities, firmware, and other files.

   For example, see the following table.

Package	What is installed
fio-util-<VSL-version>.x86_64.rpm	VSL utilities—Recommended
fio-firmware-<firmware-version>.noarch.rpm	Firmware archive—Recommended
libvsl-<version>.x86_64.rpm	Libraries needed for management tools—Recommended
fio-common-<VSL-version>.x86_64.rpm	Files required for the `init` script—Recommended
fio-sysvinit-<VSL-version>.x86_64.rpm	`Init script`—Recommended
fio-snmp-agentx-<version>.x86_64.rpm	Agentx SNMP subagent—Optional. For more information, see "Setting up SNMP for Linux (on page 27)"
fio-snmp-mib-<version>.x86_64.rpm	SNMP MIBs—Optional. For more information, see "Setting up SNMP for Linux (on page 27)"
libfio-dev-<version>.x86_64.rpm	2.x Management SDK, deprecated—Optional
libfio-doc-<version>.x86_64.rpm	2.x Management SDK, deprecated—Optional
libvsl-dev-<version>.x86_64.rpm	Current Management SDK—Optional
libvsl-doc-<version>.x86_64.rpm	Current Management SDK—Optional

4. Build your binary `rpm` from the source `rpm`. For more information, see "Building the IO Accelerator driver from source (on page 14)." Return to this step when the binary driver `rpm` is created.
5. Change to the directory where the installation packages were downloaded.
6. To install the custom-built software package, enter the following command, using the package name that you just copied or downloaded into that directory.

   ```
   rpm -Uvh iomemory-vsl-<kernel-version>-<VSL-version>.x86_64.rpm
   ```

7. To install the support files, enter the following commands:

   ```
   rpm -Uvh lib*.rpm
   rpm -Uvh fio*.rpm
   ```

   The IO Accelerator and utilities are installed to the following locations.

Package type	Installation location
IO Accelerator	`/lib/modules/<kernel-version>/extra/fio/iomemory-vsl.ko`
Utilities	`/usr/bin`
Firmware	`/usr/share/fio/firmware`

Package type	Installation location
SNMP MIB	/usr/share/fio/mib

IMPORTANT: HP IO Accelerator Management Tool 3.0 Installation
HP IO Accelerator Management Tool 3.0 is a free GUI solution for managing IO Accelerator devices. The tool is also available from the HP website (http://www.hp.com/support). Uninstall any previous versions of HP IO Accelerator Management Tool before installing the latest version. To install and use HP IO Accelerator Management Tool, download and follow the installation and user guides located in the ioSphere download folder.

When all package installations are complete, go to "Loading the IO Accelerator driver (on page 18)."

Building the IO Accelerator driver from source

The IO Accelerator driver is distributed as a source package. If a binary version of the software is not available, you must build the IO Accelerator driver from source. Use the source package that is made for your distribution. Source packages from other distributions might not work.

1. Download current IO Accelerator source and support packages from the HP website (http://www.hp.com/support).

   IMPORTANT: The exact source package to download depends on your operating system, but it is either an RPM package (for operating systems that use RPM packages) or a tar installation package (for all other operating systems).

2. Change directory to where you downloaded the source package.
3. To create a customized installation package, follow the instructions in "Building an RPM installation package (on page 14)".

Building an RPM installation package

1. Install the prerequisite files for your kernel version.

   IMPORTANT: Some of the prerequisite packages might currently be in the default OS installation. If the system is not configured to get packages over the network, you might have to mount the install CD/DVD.

   - On RHEL 5/6, you need `kernel-devel`, `kernel-headers`, `rpm-build`, GCC4, and `rsync`:

     ```
     $ yum install kernel-devel kernel-headers rpm-build gcc rsync
     ```

   CAUTION: The `yum` option might not install the correct (matching) kernel versions of `kernel-devel` or `kernel-headers` packages. The `yum` option will download the latest version. Use the following command to force `yum` to download the exact versions:

   ```
   yum install kernel-headers-`uname -r`
   kernel-devel-`uname -r` gcc rsync rpm-build
   ```

   If the exact versions are no longer available in the repository, you must manually download them from the Internet. For more information, contact HP Support (http://www.hp.com/support).

   - On SLES 10/11 you need `kernel-syms`, `make`, `rpm-build`, GCC4, and `rsync`:

     ```
     $ zypper install kernel-syms make rpm-build gcc rsync
     ```

2. Build an RPM installation package for the current kernel:

   ```
   $ rpmbuild --rebuild iomemory-vsl-<VSL-version>.src.rpm
   ```

IMPORTANT: If your kernel is a UEK, you also may need to use the `--nodeps` option.

When using a .rpm source package for a non-running kernel, run the following command:

```
$ rpmbuild --rebuild --define 'rpm_kernel_version <kernel-version>'
iomemory-vsl-<VSL-version>.src.rpm
```

The new RPM package is located in a directory that is indicated in the output from the `rpmbuild` command. To find the package, look for the **Wrote** line. In the following example, the RPM packages are located in the `/usr/src/redhat/RPMS/x86_64/` directory:

```
...
Processing files: iomemory-vsl-source-2.2.3.66-1.0.x86_64.rpm
Requires(rpmlib): rpmlib(PayloadFilesHavePrefix) <= 4.0-1
rpmlib(CompressedFileNames) <= 3.0.4-1
Obsoletes: iodrive-driver-source
Checking for unpackaged file(s): /usr/lib/rpm/check-files
/var/tmp/iomemory-vsl-2.2.3.66-root
```

Wrote:

```
/usr/src/redhat/RPMS/x86_64/iomemory-vsl-2.6.18-128.el5-2.2.3.66-1.0.x86
_64.rpm
```

Wrote:

```
/usr/src/redhat/RPMS/x86_64/iomemory-vsl-source-2.2.3.66-1.0.x86_64.rpm
```

3. Record the RPM location; you will need this information later in the installation.

   The installation packages are now created for your distribution and kernel.

4. Copy the custom-built software installation RPM package into the directory where you downloaded the installation packages.
5. Return to "Installing RPM packages ("Installing RPM packages on SUSE, RHEL, and OEL" on page 12)."

Upgrading device firmware from VSL 1.x.x or 2.x.x to 3.x.x

CAUTION: You cannot downgrade an HP IO Accelerator device firmware to an earlier version after you have upgraded the device.

CAUTION: Upgrading IO Accelerator devices that were previously configured for VSL 1.x.x or 2.x.x to work with VSL 3.x.x requires a low-level media format of the device. No user data is maintained during the media format process. Be sure to backup all data on your IO Accelerator device as instructed before upgrading the firmware.

Version 3.2.3 of the HP IO Accelerator VSL supports new features, including the latest generation of IO Accelerator architecture and improved Flashback protection. These features require the latest version of the firmware. Every IO Accelerator device in a system running 3.1.x or later must be upgraded to the latest version of the firmware.

For example, if you have a system running 2.3.1 HP IO Accelerator VSL with IO Accelerator devices previously installed, and you want to install new IO Accelerator Gen2 devices (that require the latest version of the firmware), then you will need to upgrade all of the existing devices to the latest firmware version.

Upgrade path

Depending on the current version of your HP IO Accelerator device, to preserve the internal structure of the device, you might have to perform multiple upgrades. The following path is the minimum upgrade path that you must follow. Upgrade the HP IO Accelerator VSL software on the system, and upgrade the firmware to the compatible version in the following order:

1.2.4 > 1.2.7 > 2.1.0 > 2.2.3 > 3.2.x

For VSL upgrade information for the HP IO Accelerator, see the *HP IO Accelerator Release Notes* on the HP website (http://www8.hp.com/us/en/support-drivers.html). General upgrade instructions, including the firmware update instructions, are available in the *HP IO Accelerator User Guide* for each operating system.

Overformatting not supported

The `-o overformat` option is not supported in the 3.x.x VSL software. All upgraded HP IO Accelerator devices are formatted to the maximum advertised capacity, regardless of whether the device was overformatted prior to the upgrade.

Upgrading procedure

Be sure to follow the upgrade path and make sure that all previously installed IO Accelerator devices are updated with the appropriate 3.2.3-compatible firmware.

If you plan to use IO Accelerator Gen1 devices and IO Accelerator Gen2 devices in the same host, perform this upgrade on all existing IO Accelerator Gen1 devices before installing the new IO Accelerator Gen2 devices.

1. Prepare each existing IO Accelerator device for upgrade:
   a. Backup user data on each IO Accelerator device.

   △ **CAUTION:** Upgrading IO Accelerator devices that were previously configured for VSL 1.x.x or 2.x.x to work with VSL 3.x.x requires a low-level media format of the device. No user data is maintained during the media format process. Be sure to backup all data on your IO Accelerator device as instructed before upgrading the firmware.

   Do not back up the data onto another IO Accelerator device on the same system. The backup must be to a local disk or to an externally attached volume.

   b. Run the `fio-bugreport` command-line utility and save the output. This output captures the device information for each device in the system. This device information will be useful in troubleshooting any upgrade issues. For example:

   ```
   fio-bugreport
   ```

   c. Detach the IO Accelerator devices. For example:

   ```
   fio-detach /dev/fct*
   ```

   For more information, see "`fio-detach` (on page 42)."
2. Unload the current IO Accelerator driver. For example:

   ```
   $ modprobe -r iomemory-vsl
   ```
3. Uninstall the 2.x HP IO Accelerator VSL software.
   a. To uninstall the software, specify the kernel version of the package you are uninstalling. Run the following command to find the installed packages:

   ```
   $ rpm -qa | grep -i iomemory
   ```

b. To uninstall the FSL, run a command similar to the following example. Specify the kernel version of the package you want to uninstall:

```
$ rpm -e iomemory-vsl-2.6.18-194.el5-2.2.0.82-1.0
```

c. To uninstall the utilities, run the following command:

```
$ rpm -e fio-util fio-snmp-agentx fio-common fio-firmware iomanager-gui
iomanager-jre libfio libfio-doc libfusionjni fio-sysvinit fio-smis
fio-snmp-mib libfio-deb
```

4. Install the new VSL and related packages:

a. Download the VSL binary package for your kernel and all supporting packages from the HP website (http://www.hp.com/support).

If you do not see a binary package for your kernel, see "Building the IO Accelerator driver from source (on page 14)." To see your current kernel version, run the following command:

```
uname -r
```

b. Install the VSL and utilities using the following commands:

```
rpm -Uvh
iomemory-vsl-<kernel-version>-<VSL-version>.x86_64.rpm
rpm -Uvh lib*.rpm
rpm -Uvh fio*.rpm
```

For more information, see "Installing RPM packages on SUSE, RHEL, and OEL (on page 12)."

c. Reboot the system.

5. Update the firmware on each device to the latest version using the `fio-update-iodrive` command.

△ **CAUTION:** Do not turn off the power during a firmware upgrade, because this might cause device failure. If a UPS is not in place, consider adding one to the system before performing a firmware upgrade.

Sample syntax:

```
fio-update-iodrive <iodrive_version.fff>
```

Where `<iodrive_version.fff>` is the path to the firmware archive. This command updates all of the devices to the selected firmware. If you wish to update specific devices, consult the utility reference for more options.

6. Reboot the system.

A warning that the upgraded devices are missing a lebmap, if `fio-status` is run. This is a customary warning, and the issue will be corrected in the next step.

7. Load the VSL. For example:

```
$ modprobe iomemory-vsl
```

For more information, see "Loading the IO Accelerator driver (on page 18)."

△ **CAUTION:** Use this utility with care since it deletes all user information on the IO Accelerator.

8. Using the `fio-format` command, format each device. For example:

```
fio-format <device>
```

You are prompted to confirm you want to erase all data on the device. The format might take an extended period of time, depending on the wear on the device.

9. Attach all IO Accelerator devices using the following command:

```
fio-attach /dev/fct*
```

10. Using the following command, check the status of all devices:

```
fio-status -a
```

Your IO Accelerator devices are now successfully upgraded for this version of the HP IO Accelerator. You can now install any IO Accelerator Gen2 devices.

Loading the IO Accelerator driver

1. Load the driver:

```
$ modprobe iomemory-vsl
```

The driver automatically loads at system boot. The IO Accelerator is now available to the operating system as `/dev/fiox`, where `x` is a letter.

For this command to work on SLES 10 systems, you must edit the `/etc/init.d/iomemory-vsl` file init info and change `udev` to `boot.udev`. The file should look like the following:

```
### BEGIN INIT INFO
# Provides: iomemory-vsl
# Required-Start:   boot.udev
```

On SLES systems you must allow unsupported modules for this command to work:

- SLES 11 Update 2: Modify the `/etc/modprobe.d/iomemory-vsl.conf` file, and then uncomment the appropriate line:

```
# To allow the ioMemory VSL driver to load on SLES11,
uncomment below
allow_unsupported_modules 1
```

- SLES 10 SP4: Modify the `/etc/sysconfig/hardware/config` file so the `LOAD_UNSUPPORTED_MODULES_AUTOMATICALLY` sysconfig variable is set to yes:

```
LOAD_UNSUPPORTED_MODULES_AUTOMATICALLY=yes
```

2. Confirm that the IO Accelerator device is attached:

```
fio-status
```

The output lists each drive and status (attached or unattached).

> **NOTE:** If the IO Accelerator device does not automatically attach, then check the `/etc/modprobe.d` files to see if the `auto_attach` option is turned off (set to `0`).

Controlling IO Accelerator driver loading

Control driver loading through the `init` script or through the `udev` command.

In newer Linux distributions, users can rely on the `udev` device manager to automatically find and load drivers for their installed hardware at boot time, though `udev` can be disabled and the `init` script used in nearly all cases. For older Linux distributions without this functionality, users must rely on a boot-time `init` script to load needed drivers. HP Support can provide an `init` script in `/etc/init.d/iomemory-vsl` to load the VSL driver in older RHEL4 releases and SLES10 distributions.

Using the init script

On systems where `udev` loading of the driver does not work or is disabled, the `init` script might be enabled to load the driver at boot. On some distributions, it might be enabled by default.

NOTE: The `init` script is part of the `fio-sysvinit` package, which must be installed before you can enable the `init` script.

To disable this loading of the IO Accelerator driver, enter the following command:

```
$ chkconfig --del iomemory-vsl
```

To re-enable the driver loading in the init script, enter the following command:

```
$ chkconfig --add iomemory-vsl
```

For more details, see "Using the init script (on page 19)".

Using udev

On systems that rely on `udev` to load drivers, if you want to prevent `udev` from auto-loading the IO Accelerator at boot time, users must modify a drivers option file.

To modify the drivers option file:

1. Locate and edit the `/etc/modprobe.d/iomemory-vsl.conf` file, which contains the following line:

   ```
   # blacklist iomemory-vsl
   ```

2. To disable loading, remove the `#` from the line, and then save the file.
3. Reboot Linux. The IO Accelerator will not load from `udev`.
4. To restore the udev-loading of the IO Accelerator, replace the `#` to comment out the line.

CAUTION: The version of udev on RHEL4u7/CentOS4u7 and earlier does not support the blacklist directive. Even if the driver is blacklisted as documented, `udev` will load the driver. To blacklist the driver in these versions, put the name of the driver on a separate line in the `/etc/hotplug/blacklist` file.
For example: `iomemory-vsl`

Disabling the loading on either udev or init script systems

Users can disable the loading of the IO Accelerator driver at boot time on either `udev` or `init` script systems. Disabling prevents the `auto_attach` process for diagnostic or troubleshooting purposes. To disable or enable the `auto_attach` functionality, follow the steps in "Disabling auto_attach ("Disabling auto attach" on page 38)."

Alternatively, you can prevent the driver from loading by appending the following parameter at the kernel command line of your boot loader:

```
iodrive=0
```

However, HP does not recommend this method, because it prevents the driver from functioning at all and limit the amount of troubleshooting you can perform.

Using the init script

The installation process places an `init` script in the `/etc/init.d/iomemory-vsl` file. This script uses the setting options found in the options file in `/etc/sysconfig/iomemory-vsl`. The options file must have ENABLED set (non-zero) for the `init` script to be used:

```
ENABLED=1
```

The options file contains documentation for the various settings: two of which, `MOUNTS` and `KILL_PROCS_ON_UMOUNT`, are discussed in more detail in the "Handling IO Accelerator driver unloads (on page 20)."

Mounting filesystems

The IO Accelerator driver is not loaded in the `initrd`, and, (built kernel) therefore, using the standard method for mounting filesystems (`/etc/fstab`), does not work.

To set up auto-mounting of a filesystem hosted on an IO Accelerator:

1. Add the `mnt` command to `/etc/fstab`.
2. Add the `noauto` option to `/etc/fstab`.

   For example:

   - `/dev/fcta /mnt/fioa ext3 defaults,noauto 0 0`
   - `/dev/fctb1 /mnt/iodrive ext3 defaults,noauto 0 0`

To have the `init` script mount these drives after the driver is loaded and unmount them before the driver is unloaded, add a list of mount points to the options file. For more information, see "Using module parameters (on page 20)."

For the filesystem mounts shown in the previous example, the line in the options file appears similar to the following:

```
MOUNTS="/mnt/fioa /mnt/iodrive"
```

Handling IO Accelerator driver unloads

By default, the `init` script searches for any processes holding open a mounted filesystem, kills them, and then enables the filesystem to be unmounted. This behavior is controlled by the option `KILL_PROCS_ON_UMOUNT` in the options file. If these processes are not killed, then the filesystem cannot be unmounted. This might keep the IO Accelerator from unloading cleanly, causing a significant delay on the subsequent boot.

Setting the IO Accelerator driver options

This section explains how to set IO Accelerator options.

Using module parameters

The following table describes the module parameters you can set by editing the `/usr/modprobe.d/iomemory-vsl.conf` file and changing the values.

IMPORTANT: To take effect, these changes must be completed before the IO Accelerator is loaded.

Module parameter	Default (minimum/maximum)	Description
`auto_attach`	True	Attach the device on startup.

Module parameter	Default (minimum/maximum)	Description
fio_dev_wait_timeout_secs	30	Number of seconds to wait for /dev/fio* files to show up during driver load. For systems not using udev, set this parameter to 0 to disable the timeout and avoid an unneeded pause during driver load.
force_minimal_mode	False	Force minimal mode on the device.
parallel_attach	True	Enable parallel attach of multiple devices.
preallocate_memory	No devices selected	For the selected devices, pre-allocate all memory necessary to have the drive usable as swap space.
tintr_hw_wait	0 (0, 255)	Interval (microseconds) to wait between hardware interrupts, also known as interrupt coalescing. 0 is off.
use_workqueue	3 (1 or 3)	Linux only: 3 = Use standard OS I/O elevators 0 = bypass

IMPORTANT: Except for the preallocate_memory instruction, module parameters apply to all IO Accelerator devices in the system.

One-time configuration

The IO Accelerator driver options can be set when the driver is installed on the command line of either insmod or modprobe. For example, set the auto_attach driver option to 0:

```
$ modprobe iomemory-vsl auto-attach=0
```

This option takes effect only for this load of this driver. This option is not set for subsequent calls to modprobe or insmod.

Persistent configuration

To maintain a persistent setting for an option, add the option to the /etc/modprobe.d/iomemory-vsl.conf file or a similar file. To prevent the IO Accelerator from auto-attaching, add the following line to the iomemory-vsl.conf file:

```
options iomemory-vsl auto_attach=0
```

The driver option then takes effect for every subsequent driver load, as well as on autoload of the driver during boot time.

Upgrading the firmware

After the IO Accelerator driver is loaded, ensure that the IO Accelerator device firmware is up to date by running the "fio-status (on page 45)" command-line utility.

If the output shows that the device is running in Minimal mode, download the latest firmware from the HP website (http://www.hp.com/support), and then use the HP IO Accelerator Management Tool application or the "fio-update-iodrive (on page 49)" utility to upgrade the firmware.

CAUTION: Upgrade Path

- Do not attempt to downgrade the firmware on any IO Accelerator device.
- You must follow a specific upgrade path when upgrading an IO Accelerator device.
- When installing a new IO Accelerator device along with existing devices, you must upgrade all of the existing devices to the latest available versions of the firmware before installing the new devices.
- Consult the release notes for this IO Accelerator release before upgrading IO Accelerator devices.

IMPORTANT: The IO Accelerator device might have a minimum firmware label affixed (for example, "MIN FW: XXXXXX"). This label indicates the minimum version of the firmware that is compatible with your device.

Enabling PCIe power

For PCIe IO Accelerators, if you have installed any dual IO Accelerator devices, such as the HP ioDrive2 Duo, then the device might require additional power than the minimum 25 W provided by PCIe Gen2 slots to properly function.

For instructions on enabling the device to draw additional power from the PCIe slots, see "Enabling PCIe power override (on page 36)."

Using the device as swap

To safely use the IO Accelerator as swap space, you must pass the `swap_mode=1` kernel module parameter. To pass this parameter, add the following line to the `/etc/modprobe.d/iomemory_vsl` file:

```
options iomemory_vsl preallocate_memory=1072,4997,6710,10345
```

where `1072,4997,6710,10345` are serial numbers obtained from the `fio-status` utility. Be sure to use serial numbers for the IO Accelerator modules and not the adapter.

A 4K sector size format is required for swap. This format reduces the driver memory footprint.

CAUTION: You must have 400MB of free RAM per 80GB of IO Accelerator capacity (formatted to 4KB block size) to enable the IO Accelerator with pre-allocation enabled for swap. Attaching an IO Accelerator with pre-allocation enabled and insufficient RAM might result in the loss of user processes and system instability.

IMPORTANT: During the loading of the IO Accelerator, the `preallocate_memory` parameter is recognized and the memory allocates when the specific device is attached.

Using the Logical Volume Manager

If you add the IO Accelerator as a supported type, the LVM volume group management application handles mass storage devices like the IO Accelerator. To use the LVM:

1. Locate the `/etc/lvm/lvm.conf` configuration file.

2. Edit the file to add an entry similar to the following:

   ```
   types = [ "fio", 16 ]
   ```

The parameter "16" represents the maximum number of partitions supported by the drive. For the IO Accelerator, this can be any number from 1 upwards. Do not set this parameter to 0.

IMPORTANT: Do not run the `udev` command to load the IO Accelerator driver while LVM or MD is active. The `init` script will disconnect from the LVM volumes and MD devices before disconnecting from the Memory device.

NOTE: For the IO Accelerator, HP recommends 16 for the partition setting.

Configuring RAID

Where possible, you can configure two or more IO Accelerators into a RAID array, using standard Linux procedures.

NOTE: If you are using RAID1 mirroring and one device fails, enter the `fio-format` command on the replacement device (not the existing good device) before rebuilding the RAID. HP recommends that you do not use a RAID5 configuration.

RAID 0

To create a striped set where `fioa` and `fiob` are the two IO Accelerators you want to stripe, enter the following command:

```
$ mdadm --create /dev/md0 --chunk=256 --level=0 --raid-devices=2 /dev/fioa
/dev/fiob
```

Making the Array Persistent (Available after restart)

IMPORTANT: On some versions of Linux, the configuration file is in the `/etc/mdadm/mdadm.conf` directory, not the `/etc/mdadm.conf` directory.

1. Inspect the `/etc/mdadm.conf` file.
2. If one or more lines declare the devices to inspect, then make sure those lines specify `partitions` as an option.
3. If it does not, add a new `DEVICE` line to the file specifying `partitions`.

   For example:

   ```
   DEVICE partitions
   ```

   Add a device specifier for the `fio ioMemory` devices:

   ```
   DEVICE /dev/fio*
   ```

4. Verify if updates are required to `/etc/mdadm.conf` by entering the following command:

   ```
   $ mdadm --detail --scan
   ```

5. Compare the output of this command to what currently exists in `mdadm.conf`, and then add any sections required to:

   `/etc/mdadm.conf`.

IMPORTANT: For example, if the array consists of two devices, the command output will not display three lines in the `mdadm.conf` file–
one line for the array, and a line for each device.
Ensure that those lines are added to the `mdadm.conf` file so it matches the output of the command.

For more details, see the `mdadm` and `mdadm.conf` manpages for your distribution. With these changes, on most systems the RAID 0 array will be created automatically upon restart.

6. If you cannot access `/dev/md0` after restart, run the following command:

   ```
   $ mdadm --assemble --scan
   ```

   You might also want to disable `udev` loading of the IO Accelerator driver, if needed, and then use the `init` script provided for driver loading. For more information, see "Using the Init Script (on page 19)."

Making an array persistent

NOTE: On some versions of Linux, the configuration file is located in the `/etc/mdadm/mdadm.conf`, not `/etc/mdadm.conf` directory.

Inspect `/etc/mdadm.conf`. If one or more lines declare the devices to inspect, be sure one of those lines specifies "partitions" as an option by adding a new `DEVICE` line to the file specifying the "partitions" option:

```
DEVICE partitions
```

Add a device specifier for the `fio` ioMemory devices:

```
DEVICE /dev/fio*
```

To verify whether any updates are needed to `/etc/mdadm.conf`, issue the following command:

```
$ mdadm --detail --scan
```

Compare the output of this command to what currently exists in `mdadm.conf` and add any needed sections to `/etc/mdadm.conf`.

IMPORTANT: For example, if the array consists of two devices, the command output will not display three lines in the `mdadm.conf` file–
one line for the array, and a line for each device.
Ensure that those lines are added to the `mdadm.conf` file so it matches the output of the command.

For more details, see the `mdadm and mdadm.conf` man pages for your distribution.

With these changes, the RAID 0 array will be created automatically upon restart on most systems. If you have problems accessing `/dev/md0` after restart, run the following command:

```
$ mdadm --assemble --scan
```

If needed, you can also disable `udev` loading of the IO Accelerator driver. To load the driver, use the `init` script.

IMPORTANT: In SLES 11, to be sure these services are run on boot, you might have to run the following commands:

- `chkconfig boot.md on`
- `chkconfig mdadmd on`

RAID 1

Create a mirrored set by using the `fioa` and `fiob` IO Accelerators:

```
$ mdadm --create /dev/md0 --level=1 --raid-devices=2 /dev/fioa /dev/fiob
```

To view your specific names, use the `fio-status` command.

RAID 10

Create a striped, mirrored array by using four IO Accelerators (`fioa`, `fiob`, `fioc`, and `fiod`):

```
$ mdadm --create /dev/md0 -v --chunk=256 --level=raid10 --raiddevices=4
/dev/fioa /dev/fiob /dev/fioc /dev/fiod
```

View the results:

```
fio-status
```

Building a RAID 10 across multiple devices

In a RAID 10 configuration, sets of two disks are mirrored, and then those mirrors are striped. When setting up a RAID 10 configuration across multiple IO Accelerator Duos, HP recommends that you make sure that no mirror resides solely on the two IO Accelerator modules that comprise an IO Accelerator Duo.

To lay the data out:

1. Use the `--layout=n2` option when creating the RAID 10 configuration. This is the default.
2. Ensure that no two IO Accelerator modules from the same duo are listed side-by-side.

The following sample code illustrates HP recommended configurations.

CAUTION: You must list the `fiox` devices in the correct sequence.

IMPORTANT: When the IO Accelerator devices have been formatted with the `fio-format` utility, use the following commands.

```
# 2 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 4 /dev/md0 \
/dev/fioa /dev/fioc \
/dev/fiob /dev/fiod
# Mirror groups are: fioa,fioc and fiob,fiod
# 3 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 6 /dev/md0 \
/dev/fioa /dev/fiod \
/dev/fioc /dev/fiof \
/dev/fioe /dev/fiob
# 4 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 8 /dev/md0 \
/dev/fioa /dev/fiod \
```

```
/dev/fioc /dev/fiof \
/dev/fioe /dev/fioh \
/dev/fiog /dev/fiob
# 8 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 16 /dev/md0 \
/dev/fioa /dev/fiod \
/dev/fioc /dev/fiof \
/dev/fioe /dev/fioh \
/dev/fiog /dev/fioj \
/dev/fioi /dev/fiol \
/dev/fiok /dev/fion \
/dev/fiom /dev/fiop \
/dev/fioo /dev/fiob
```

Understanding Discard (TRIM) support

Discard (also known as TRIM) is enabled by default in this version of the IO Accelerator driver.

Discard addresses a unique issue to solid-state storage. When a user deletes a file, the device does not recognize that it can reclaim the space. Instead, the device assumes the data is valid.

Discard is a feature on newer filesystem releases. It informs the device of logical sectors that no longer contain valid user data. This enables the wear-leveling software to reclaim that space (as reserve) to handle future write operations.

Discard TRIM on Linux

Discard is enabled by default. For Discard to be implemented, the Linux distribution must support this feature, and Discard must be enabled.

Under Linux, discards are not limited to being created by the filesystem; discard requests can also be generated directly from userspace applications using the kernel `discard ioctl`.

CAUTION: A known issue is that ext4 in Kernel.org 2.6.33 or earlier might silently corrupt data when Discard is enabled.
The issue has been fixed in many kernels provided by distribution vendors. Check with your kernel provider to ensure that your kernel properly supports Discard. For more information, see the release notes for this version of the driver.

IMPORTANT: Currently, MD and LVM do not pass discards to underlying devices in Linux. Therefore, any ioDrive device that is part of an MD or LVM array will not receive discards sent by the filesystem.

The LVM release included in Red Hat 6.1 supports passing discards for several targets, but not all (RHEL 6.1documentation (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/newmds-ssdtuning.html)). For more information, see your distribution documents.

Setting up SNMP for Linux

SNMP details for Linux

The `fio-snmp-agentx` SNMP agent is an RFC 2741-compliant AgentX subagent. It can work with any RFC-compliant SNMP agent, such as Net-SNMP.

The master SNMP agent defers queries to `fio-snmp-agentx` for supported MIBs.

Files and directories

File	Description
/usr/share/fio/mib/cpqIODrive.mib	HP MIB for the IO Accelerator
/opt/fio/etc/snmp/fio-snmp-agent.conf*	HP IO Accelerator SNMP configuration
/etc/snmp/snmpd.conf	Master SNMP configuration file
/var/log/fio/fio-snmp-agent.log	Log file for the subagent
/usr/share/doc/fio/fio-snmp-agentx/conf	Sample configuration and `testini` files
`hp-ioaccel-snmp-agent-<version>.x86_64.rpm`	RPM that configures the `/etc/snmpd.conf` file and the `fio-snmp-agentx.conf` file

* This file might be a symbolic link to `/etc/snmp/fio-snmp-agentx.conf`.

SNMP master agent

The `fio-snmp-agentx`, provided in the `fio-util` package, requires an already-installed SNMP master agent. The SNMP master agent must support and be configured for AgentX connections. For more information, see the "*AgentX Protocol Version 1* (*http://www.ietf.org/rfc/rfc2741.txt*)."

The `fio-snmp-agentx` is tested and verified with Net-SNMP, which is the typical SNMP agent provided with most Linux distributions. Several agents are available that support this functionality. The following sections describe Net-SNMP.

Launching the SNMP master agent

Install the `Net-SNMP` package using the package manager for your version of Linux.

Red Hat

To install `Net-SNMP` on Red Hat, use the following command:

```
yum install net-snmp rsync
```

Other Linux versions

To install the `Net-SNMP` package on your Linux distribution, use the standard system package manger. The `fio-snmp-mib` package places MIB files in `/usr/share/fio/mib` directory.

Configuring the SNMP master agent

You can configure the `Net-SNMP` master agent daemon to set the network communications parameters, security, and other options by using the `snmpd.conf` text file. The location of this file is system-dependent and customarily found in the following directory:

```
/etc/snmp or /usr/share/snmp.
```

A simple `snmpd` configuration file might include the following:

```
# set standard SNMP variables
syslocation "Data room, third rack"
syscontact itguy@example.com
# required to enable the AgentX protocol
master agentx
agentxsocket tcp:localhost:16101
#set the port that the agent listens on (defaults to 161)
agentaddress 161
# simple access control (some form of access control is required)
rocommunity public
```

Running the Master Agent

After you install and configure the master agent, you must start or restart the `snmpd` daemon for the new parameters to take effect. You can run `snmpd` from the installed location (customarily located in the `/usr/sbin` directory. For options, see the `snmpd` manpage). To run properly, the `snmpd` daemon typically needs `root` privileges. You can also use the `snmpd` startup script in the

`/etc/init.d` or `/etc/rc.d/init.d` directory. For additional security, use the advanced SNMPv3 access control instead of the `rocommunity` and `rwcommunity` access control directives as outlined in the relevant manpage.

SNMP agentX subagent

△ **CAUTION:** The SNMP agent requires the `libvsl RPM` package. Install this agent as part of the IO Accelerator installation. For more information, contact HP Support (http://www.hp.com/support).

Installing the SNMP subagent

1. The SNMP package is part of the overall software package previously installed. Download the IO Accelerator Linux packages from the HP website (http://www.hp.com/support).

2. Install the package using your operating systems package manager. For example, on Red Hat:

   ```
   rpm -Uvh fio-snmp-*.rpm
   ```

   The SNMP package places the MIB files in the `/usr/share/fio/mib` directory.

Running and configuring the SNMP subagent

An RPM to configure the SNMP files is available on the HP website (http://www.hp.com/go/support). The RPM is named hp-ioaccel-snmp-agent1-1.x86_64.rpm.

To manually set up SNMP:

1. Configure the subagent by creating a `fio-snmp-agentx.conf` file.
2. Store the `fio-snmp-agentx.conf` file in the `/opt/fio/etc/snmp` directory.
3. Set the agent network parameters in this file similar to the following:

   ```
   # required to enable the AgentX protocol
   agentxsocket tcp:localhost:16101
   ```

This must match the AgentX network parameters in the `snmpd.conf` file for the master agent. For further AgentX configuration information, consult the manpages or contact HP Support (http://www.hp.com/support).

The `fio-snmp-agentx` startup script will launch automatically at boot time after the installation and configuration is complete.

Manually running the SNMP subagent

1. After the SNMP master agent is started, start the subagent:

   ```
   /usr/bin/fio-snmp-agentx
   ```

   This command launches the subagent using the Net-SNMP `fio-snmp-agentx.conf` configuration file. This file must reside in one of the `/opt/fio/etc/snmp` directories.
2. View the IO Accelerator management information by using an SNMP MIB browser, HP system management homepage, or by using a network management system accessing `cpqIODDrv.mib`, located in the `/usr/share/fio/mib` directory.

Subagent log file

The HP IO Accelerator subagent can maintain a log file regarding its own activities. This file is separate from the MIB because it includes entries on the subagent communications with the master agent including any errors or intermittent issues.

To have the subagent maintain this log file, include the `-l` parameter and a path to the log file as part of the command in running the subagent.

The default log file is `/var/log/fio/fio-snmp-agentx.log`.

For example, running the following command keeps the subagent log files as `subagent.log` in the `/usr/snmp` directory:

```
fio-snmp-agentx -l /usr/snmp/subagent.log
```

The SNMP subagent is now ready to monitor your device.

Using the SNMP sample config files

When you install SNMP, the following sample `config` files are available:

- `/usr/share/doc/fio-snmp-agentx/conf/snmpd.conf.hp/` (master agent)
- `/usr/share/doc/fio-snmp-agentx/conf/fio-snmp-agentx.conf.hp/` (subagent)

To customize and use the sample `config` files:

1. Rename the `snmpd.conf` and the `fio-snmp-agentx.conf` files. For example, `snmpd.orig.conf` and `fiosnmp-agentx-orig.conf`.

   The `snmpd.conf` file is located in `/etc/snmp` or the `/usr/share/snmp` directories. The `fio-snmp-agentx.conf` file is located in the `/opt/fio/etc/snmp` directory.

   NOTE: If the `hp-ioaccel-snmp-agent.rpm` file is installed, the link is symbolic.

2. From the `/usr/share/doc/fio-snmp-agentx/conf/` directory, copy the sample `snmpd.conf.hp` and the sample `fio-snmp-agentx.conf.hp` files to the appropriate directories.
3. Rename the files without the `.hp` extension. For example, rename the `snmp.conf.hp` file to `snmp.conf` and rename `fio-snmp-agentx.conf.hp` to `fio-snmp-agentx.conf`.
4. Edit the sample files, and then save the changes as `snmpd.conf` and `fio-snmp-agentx.conf`.

Enabling SNMP test mode

When the SNMP subagent runs, it reads the `fio-snmp-agentx` config file:

```
######################################################################
# Example config file for fio-snmp-agentx SNMP AgentX subagent.
#
agentxsocket tcp:localhost:16101
# test_mode_enabled
# set to 1, true or yes to enable 0, false or no to disable (default: false)
test_mode_enabled true
# traps_enabled
traps_enabled true
# testmode_file
# name of test mode file (default: testmode.ini)
testmode_file testmode.ini
# update_delay
# delay between agent polling requests in milliseconds (default: 250)
update_delay 100
# mib_select
# set to cpq for CPQIODRV-MIB
mib_select cpq
######################################################################
```

The conditions for test mode include the following:

- If the Admin has set the `test_mode_enabled` parameter from TRUE to FALSE, then the SNMP subagent does not attempt to run test mode but it continues processing data as usual from the IO Accelerator driver, storing the data in the MIB.
- If the CONF file says that `test_mode_enabled` is TRUE, then the SNMP subagent first reads the `testmode_file` line to locate the `testmode.ini` file. Next, the subagent reads this file.
- If the `testmode.ini` file shows the test mode is set to ON, then it engages the test mode.
- If test mode is ON, then the SNMP subagent reads the next line, `TestModeIndex`, to identify which IO Accelerator to test. The number in this parameter is the PCIe device number shown using `fio-status` such as:

  `PCI:01:00.0`

  The first two numerals identify the PCIe bus number (in this case, `01`). This bus number is reported in hexadecimal, whereas the `TestModeIndex` in the `testmode.ini` file must be specified in decimal. The converted number must be entered into `testmode.ini` file. The `TestModeIndex` must be a valid bus number of an IO Accelerator installed in the system.

  The SNMP subagent now replaces any existing IO Accelerator driver data it might have for the IO Accelerator specified by `TestModeIndex` with any populated fields in the list of parameters. If a field is not populated, then the subagent retains the existing data and reports it to the MIB. If the field has a value, then the subagent replaces that data, and reports it to the MIB.

  The subagent continues in test mode until the `.ini` file parameter is set to OFF. The test mode information is described in the `testmode.ini` file. A sample.ini file is located in the `/usr/share/doc/fio/fio-snmp-agentx/conf` directory:

```
# SNMP Test Mode sample file.

# These values may be used to test the SNMP subsystem when it is in test mode.

[SNMP Agent Test Mode]

TestMode = off

TestModeIndex = 0

# InfoState: Note that the following states may change, but current definitions
are:

# 0 = unknown

# 1 = detached

# 2 = attached

# 3 = minimal mode

# 4 = error

# 5 = detaching

# 6 = attaching

# 7 = scanning

# 8 = formatting

# 9 = updating firmware

# 10 = attach

# 11 = detach

# 12 = format
```

```
# 13 = update
InfoState = 2
InfoInternalTemp = 45
InfoAmbientTemp = 35
InfoWearoutIndicator = 2 ; 2=normal, 1=device is wearing out.
InfoWritableIndicator = 2 ; 2=normal, 1=non-writable, 0=write-reduced, 3=unknown
InfoFlashbackIndicator = 2 ; 2=normal, 1=flashback protection degraded.
ExtnTotalPhysCapacityU = 23
ExtnTotalPhysCapacityL = 215752192
ExtnUsablePhysCapacityU = 21
ExtnUsablePhysCapacityL = 7852192
ExtnUsedPhysCapacityU = 4
ExtnUsedPhysCapacityL = 782330816
ExtnTotalLogCapacityU = 18
ExtnTotalLogCapacityL = 2690588672
ExtnAvailLogCapacityU = 14
ExtnAvailLogCapacityL = 3870457856
ExtnBytesReadU = 18
ExtnBytesReadL = 3690588672
ExtnBytesWrittenU = 4
ExtnBytesWrittenL = 2578550816
InfoHealthPercentage = 95
InfoMinimalModeReason = 7 ; 0=unknown, 1=fw out of date, 2=low power, ; 3=dual
plane failure, 5=internal, 6=card limit, ; 7=not in minimal mode, 8=unsupported
OS,; 9=low memory
InfoReducedWriteReason = 0 ; 0=none, 1=user requested, 2=no md blocks, ; 3=no
memory, 4=failed die, 5=wearout, ; 6=adapter power, 7=internal, 8=power limit
InfoMilliVolts = 12000
InfoMilliVoltsPeak = 12100
InfoMilliVoltsMin = 11900
InfoMilliWatts = 6000
InfoMilliWattsPeak = 15000
InfoMilliAmps = 500
InfoMilliAmpsPeak = 1000
InfoAdapterExtPowerPresent = 1 ; 1=present, 2=absent
InfoPowerlossProtectDisabled = 2 ; 1=powerloss protection available but disabled
; 2=any other powerloss protection condition
```

Troubleshooting SNMP

For SMH issues, ensure you have installed the latest web templates available for the IO Accelerator from the HP website (http://www.hp.com/go/support). If a new PSP has been loaded, the latest IO Accelerator templates should be installed after the PSP is installed.

If the IO Accelerator is not viewable on the SMH, the port selected might already be in use.

To verify whether the port is already in use:

1. After the SNMPD service is restarted, verify that SNMP is running by entering the `ps –ef |grep snmp` command. The SNMP daemon `snmpd` must be running. If it is not running, start the SNMP service.
2. View the system log messages verifying whether any service failed. In most instances, all system errors get logged in the `/var/log/messages` file.
3. Run the `find /var/log/m* | xargs grep 16101 command`, where 16101 is the port number.

   The command might result in the following message related to socket Error: `Couldn't open a master agentx socket to listen on (tcp:localhost:16101): Unknown host (tcp:localhost:16101) (Permission denied).`
4. If this message appears, then find another available port. For example, assume 7052 is a free port. Use the `netstat –a | grep 7052` command. If it returns nothing, edit both the agent file `snmpd.conf` and the subagent file `fio-snmp-agentx.conf` to use this free port.

Supported SNMP MIB fields

SNMP MIB	SNMP MIB
cpqIoDrvMibRevMajor	cpqIoDrvInfoAdapterType
cpqIoDrvMibRevMinor	cpqIoDrvInfoAdapterPort
cpqIoDrvMIBCondition	cpqIoDrvInfoAdapterSerialNumber
cpqIoDrvInfoIndex	cpqIoDrvInfoAdapterExtPowerPresent
cpqIoDrvDimmInfoStatus	cpqIoDrvInfoPowerlossProtectDisabled
cpqIoDrvInfoName	cpqIoDrvInfoInternalTempHigh
cpqIoDrvInfoSerialNumber	cpqIoDrvInfoAmbientTemp
cpqIoDrvInfoPartNumber	cpqIoDrvInfoPCIBandwidthCompatibility
cpqIoDrvInfoSubVendorPartNumber	cpqIoDrvInfoPCIPowerCompatibility
cpqIoDrvInfoSparePartNumber	cpqIoDrvInfoActualGoverningLevel
cpqIoDrvInfoAssemblyNumber	cpqIoDrvInfoLifespanGoverningLevel
cpqIoDrvInfoFirmwareVersion	cpqIoDrvInfoPowerGoverningLevel
cpqIoDrvInfoDriverVersion	cpqIoDrvInfoThermalGoverningLevel
cpqIoDrvInfoUID	cpqIoDrvInfoLifespanGoverningEnabled
cpqIoDrvInfoState	cpqIoDrvInfoLifespanGoverningTgtDate
cpqIoDrvInfoClientDeviceName	cpqIoDrvExtnIndex
cpqIoDrvInfoBeacon	cpqIoDrvExtnTotalPhysCapacityU
cpqIoDrvInfoPCIAddress	cpqIoDrvExtnTotalPhysCapacityL
cpqIoDrvInfoPCIDeviceID	cpqIoDrvExtnTotalLogCapacityU
cpqIoDrvInfoPCISubdeviceID	cpqIoDrvExtnTotalLogCapacityL
cpqIoDrvInfoPCIVendorID	cpqIoDrvExtnBytesReadU
cpqIoDrvInfoPCISubvendorID	cpqIoDrvExtnBytesReadL
cpqIoDrvInfoPCISlot	cpqIoDrvExtnBytesWrittenU
cpqIoDrvInfoWearoutIndicator	cpqIoDrvExtnBytesWrittenL

SNMP MIB	SNMP MIB
cpqIoDrvInfoFlashbackIndicator	cpqIoDrvExtnFormattedBlockSize
cpqIoDrvInfoWritableIndicator	cpqIoDrvExtnCurrentRAMUsageU
cpqIoDrvInfoInternalTemp	cpqIoDrvExtnCurrentRAMUsageL
cpqIoDrvInfoHealthPercentage	cpqIoDrvExtnPeakRAMUsageU
cpqIoDrvInfoMinimalModeReason	cpqIoDrvExtnPeakRAMUsageL
cpqIoDrvInfoReducedWriteReason	cpqIoDrvWearoutTrap
cpqIoDrvInfoMilliVolts	cpqIoDrvNonWritableTrap
cpqIoDrvInfoMilliVoltsPeak	cpqIoDrvFlashbackTrap
cpqIoDrvInfoMilliVoltsMin	cpqIoDrvTempHighTrap
cpqIoDrvInfoMilliWatts	cpqIoDrvTempOkTrap
cpqIoDrvInfoMilliWattsPeak	cpqIoDrvErrorTrap
cpqIoDrvInfoMilliAmps	cpqIoDrvPowerlossProtectTrap
cpqIoDrvInfoMilliAmpsPeak	

Maintenance

Maintenance tools

The IO Accelerator includes software utilities for maintaining the device. You can also install SNMP as a monitoring option.

The following are the most common tasks for maintaining your IO Accelerator. You can also use the IO Accelerator Management Tool application to perform firmware upgrades. For more information, see the *HP IO Accelerator Management Tool User Guide*.

Device LED indicators

The IO Accelerator device includes three LEDs showing drive activity or error conditions.

HP IO Accelerator Management Tool

The HP IO Accelerator Tool is a GUI solution for managing IO Accelerator devices. The GUI is available from HP Support *(*http://www.hp.com/support*)*.

The HP IO Accelerator Management Tool can perform:

- Firmware upgrades
- Low-level formatting
- Attach and detach actions
- Device status and performance information

Command-line utilities

Several command-line utilities are included in the installation packages for managing your IO Accelerator device:

- `fio-attach`
- `fio-beacon`
- `fio-bugreport`
- `fio-detach`
- `fio-format`
- `fio-pci-check`
- `fio-snmp-agentx`
- `fio-status`
- `fio-sure-erase`

- `fio-update-iodrive`

For more information, see "Utilities reference (on page 40)."

Enabling PCIe power override

For PCIe IO Accelerators, if you have installed any dual IO Accelerator devices, such as the HP ioDrive2 Duo, then the device might require additional power than the minimum 25 W provided by PCIe Gen2 slots to properly function. Even if additional power is not required for your device, all dual IO Accelerator devices that receive additional power might benefit with improved performance.

HP ioDrive2 Duo devices must have additional power to properly function. For more information on which devices require additional power, see the *HP PCIe IO Accelerator for ProLiant Servers Installation Guide*.

Additional power can be provided in two ways:

- **External power cable**—This cable ships with all dual ioMemory devices. For information on installing this cable, see the *HP PCIe IO Accelerator for ProLiant Servers Installation Guide*.

  NOTE: When a power cable is used, all of the power is drawn from the cable and no power is drawn from the PCIe slot.

- **Enabling full slot power draw**—Some PCIe slots provide additional power (often up to 75 W of power). If your PCIe slot is rated to provide at least 55 W, then you can allow the device to draw full power from the PCIe slot by setting a VSL module parameter. For more information on enabling this override parameter, see "Enabling the override parameter (on page 37)."

  △ **CAUTION:** If the PCIe slot is not capable of providing the needed amount of power, then enabling full power draw from the PCIe slot might result in malfunction or even damage to server hardware. The user is responsible for any damage to equipment due to improper use of the override parameter. HP expressly disclaims any liability for damage arising from improper use of the override parameter. To confirm the power limits and capabilities of each slot, as well as the entire system, contact the server manufacturer. For information about using the override parameter, contact HP Customer Support.

  NOTE: The override parameter overrides the setting that prevents devices from drawing more than 25 W from the PCIe slot. The parameter is enabled by device using the device serial numbers. Once the setting is overridden, each device might draw up to the full 55 W needed for peak performance.

  Before you enable the override parameter, ensure that each PCIe slot is rated to provide enough power for all slots, devices, and server accessories. To determine the power slot limits, consult the server documentation, BIOS interface, setup utility, or use the `fio-pci-check` command.

Important considerations

- If you are installing more than one dual IO Accelerator device and enabling the override parameter for each device, be sure the motherboard is rated to provide 55W power to each slot used. For example, some motherboards safely provide up to 75W to any one slot, but run into power constraints when multiple slots are used to provide that much power. Installing multiple devices in this situation might result in server hardware damage. Consult with the manufacturer to determine the total PCIe slot power available.
- The override parameter persists in the system and enables full power draw on an enabled device even if the device is removed and then placed in a different slot within the same system. If the device is placed

in a slot that is not rated to provide 55W of power, your server hardware could experience a power drag.

- The override parameter is a setting for the IO Accelerator VSL software by server and is not stored in the device. When moved to a new server, the device defaults to the 25 W power limit until an external power cable is added or the override parameter is enabled for that device in the new server. To determine the total PCIe slot power available for the new server, consult the manufacturer.

Enabling the override parameter

1. Use one of the following methods to determine the serial number of each device to be installed in a compatible slot:
   - Enter the `fio-status` command:

     Sample output:

     ```
     fio-status

     ...
     Adapter: Dual Controller Adapter
     Fusion-io ioDrive2 DUO 2.41TB, Product Number:F01-001-2T41-CS-0001, FIO
     SN:1149D0969
     External Power: NOT connected
     PCIe Power limit threshold: 24.75W
     Connected ioMemory modules:
     fct2: SN:1149D0969-1121
     fct3: SN:1149D0969-1111
     ```

     In this example, `1149D0969` is the adapter serial number.

     If you have multiple IO Accelerator devices installed on your system, use the fio-beacon command to verify where each device is physically located.
   - Inspect the adapter serial number labels on the IO Accelerator devices to determine the serial numbers. However, HP recommends confirming that each serial number is an adapter serial number by running the fio-status command. The adapter serial number label resides on the back of all HP ioDrive Duo devices and HP ioDrive2 Duo devices. On ioDrive Duos, the serial number is located on the PCB component that is attached to the PCIe connector.
2. To set the module parameter, edit the `/usr/modprove.d/iomemory-vsl.conf` file, and then change the value for the `external_power_override` parameter. For example:

   ```
   options iomemory-vsl external_power_override=<value>
   ```

   Where the <value> for this parameter is a comma-separated list of adapter serial numbers. For example:

   1149D0969,1159E0972,24589
3. To enforce any parameter changes, you must reboot or unload and then load the drivers.

Common maintenance tasks

The following are the most common tasks for maintaining your IO Accelerator device using command-line utilities.

IMPORTANT: All commands require administrator privileges. To run the commands, log in as `root` or use `sudo`.

IMPORTANT: If you came to this section from the Software Installation section, return to that section after you uninstall previous versions of the driver and utilities.

Unloading the IO Accelerator driver

Unload the IO Accelerator driver:

```
$ modprobe -r iomemory_vsl
```

Uninstalling the IO Accelerator driver RPM package

Versions 1.2.x

Remove prior versions of the IO Accelerator software:

```
$ rpm -e iodrive-driver
```

Versions 2.x.x

With version 2.x.x of the IO Accelerator, you must specify the kernel version of the package that you are uninstalling.

Find the installed driver packages:

```
$ rpm -qa | grep -i iomemory
```

Sample output:

```
iomemory-vsl-2.6.18-194.el5-2.2.2.82-1.0
```

For example, to uninstall the IO Accelerator (specifying the kernel version of the driver you wish to uninstall):

```
$ rpm -e iomemory-vsl-2.6.18-194.el5-2.2.0.82-1.0
```

Uninstalling the IO Accelerator Utilities, IO Accelerator Management Tool 2.x, and other support packages

Uninstall the support RPM packages (adding or removing package names as needed):

```
$ rpm -e fio-util fio-snmp-agentx fio-common fio-firmware iomanager-gui
iomanager-jre libfio libfio-doc libfusionjni fio-sysvinit fio-smis fio-snmp-mib
libfio-dev
```

To uninstall the support DEB packages, (adding or removing package names as needed):

```
$ dpkg -r fio-util fio-snmp-agentx fio-common fio-firmware iomanager-gui
iomanager-jre libfio libfio-doc libfusionjni fio-sysvinit fio-smis fio-snmp-mib
libfio-dev
```

Disabling auto attach

When the IO Accelerator driver is installed, it is configured to automatically attach any devices when the driver is loaded. When necessary, disable the `auto_attach` feature.

To disable `auto_attach` using the Linux `init` script:

1. Edit the following file:

   `/etc/modprobe.d/iomemory-vsl.conf`
2. Add the following line to the file:

   `options iomemory_vsl auto_attach=0`
3. Save the file.
4. To re-enable `auto_attach`, edit the file by either removing the line added in step 2, or edit the line as in the following:

   `options iomemory-vsl auto_attach=1`

Unmanaged shutdown issues

If you experience an unmanaged shutdown, then the IO Accelerator performs a consistency check during the reboot. The reboot might take several minutes or more to complete and is indicated by a progress percentage during the startup.

Although data written to the IO Accelerator device is not lost due to unmanaged shutdowns, important data structures might not have been properly committed to the device. This consistency check repairs these data structures.

Disabling the driver

By default, the driver automatically loads when the operating system starts. You can disable the IO Accelerator auto-load for diagnostic or troubleshooting purposes.

To disable the IO Accelerator driver auto-load:

1. Append the following parameter to the kernel command line of your boot loader:

   `iodrive=0`

   The IO Accelerator does not load, so the device cannot be available to users.

   IMPORTANT: Disable the IO Accelerator to keep it from loading, or move it out of the `/lib/modules/<kernel_version>` directory.
2. Proceed with troubleshooting to correct the problem. If outdated firmware is the problem, use `iodrive=1` to place the IO Accelerator in minimal mode. Use `fio-update-iodrive` or HP IO Accelerator Management Tool to update the firmware.

To reenable the IO Accelerator to the system, use the `fio-attach` utility or HP IO Accelerator Management Tool.

Utilities

Utilities reference

The IO Accelerator installation packages include various command line utilities, installed by default to the `/usr/bin` file. These utilities provide a number of useful manners to access, test, and manipulate your device.

Utility	Purpose
fio-attach	Makes an IO Accelerator available to the OS
fio-beacon	Lights the IO Accelerator external LEDs
fio-bugreport	Prepares a detailed report for use in troubleshooting issues
fio-detach	Temporarily removes an IO Accelerator from OS access
fio-format	Used to perform a low-level format of an IO Accelerator
fio-pci-check	Searches for errors on the PCIe bus tree, specifically for PCIe IO Accelerators
fio-snmp-agentx	SNMP subagent that implements the SNMP FUSION-IODRV-MIB for the IO Accelerator
fio-status	Displays information about the device
fio-sure-erase	Clears or purges data from the device
fio-update-iodrive	Updates the IO Accelerator firmware

NOTE: All utilities have –h (Help) and –v (Version) options.

fio-attach

Description

Attaches the IO Accelerator device and makes it available to the operating system. This creates a block device in `/dev` named fiox (where x is a, b, c, and so on). Then, you might partition or format the IO Accelerator device, or set it up as part of a RAID array. The command displays a progress bar and percentage as it operates.

NOTE: In most cases, the IO Accelerator automatically attaches the device on load and does a scan. You only have to run `fio-attach` if you ran `fio-detach` or if you set the IO Accelerator `auto_attach` parameter to `0`.

Syntax

```
fio-attach <device> [options]
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the card number: `0`, `1`, `2`, and so on. For example, `/dev/fct0` indicates the first IO Accelerator device installed on the system.

You can specify multiple IO Accelerator devices. For example, `/dev/fct1 /dev/fct2` indicates the second and third IO Accelerator devices installed on the system. You can also use a wildcard to indicate all IO Accelerator devices on the system.

For example, `/dev/fct*`

Option	Description
`-c`	Attach only if clean.
`-q`	Quiet: disables the display of the progress bar and percentage.

fio-beacon

Description

The `fio-beacon` utility enables all three LEDs on, to identify the specified IO Accelerator device.

IMPORTANT: This utility enables the LEDs on, unless you select the `-0` option.

Syntax

```
fio-beacon <device> [options]
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the card number: `0`, `1`, `2`, and so on. For example, `/dev/fct0` indicates the first IO Accelerator device installed on the system.

Options	Description
`-0`	Off: Turns off the three LEDs
`-1`	On (default): Enables the three LEDs on
`-p`	Prints the PCI bus ID of the device at `<device>` to standard output. Usage and error information may be written to standard output rather than to standard error.

fio-bugreport

Description

Prepares a detailed report of the device for use in troubleshooting problems. The results are saved in the `/tmp` directory in a file that indicates the date and time the utility was run.

Example

```
/tmp/fio-bugreport-20100121.173256-sdv9ko.tar.bz2
```

Syntax

```
fio-bugreport
```

NOTE: If the utility recommends that you contact Fusion-io support, disregard that message and contact HP support (http://www.hp.com/support) instead.

Sample output

```
-bash-3.2# fio-bugreport /tmp/fio-bugreport-20090921.173256-sdv9ko ~
Collecting fio-status -a
```

```
Collecting fio-status
Collecting fio-pci-check
Collecting fio-pci-check -v
Collecting fio-read-lebmap /dev/fct0
Collecting fio-read-lebmap -x /dev/stdout/dev/fct0
Collecting fio-read-lebmap -t /dev/fct0
Collecting fio-get-erase-count/dev/fct0
Collecting fio-get-erase-count -b /dev/fct0
Collecting lspci
Collecting lspci -vvvvv
Collecting lspci -tv
Collecting messages file(s)
Collecting procfusion file(s)
Collecting lsmod
Collecting uname -a
Collecting hostname
Collecting sar -r
Collecting sar
Collecting sar -A
Collecting syslog file(s)
Collecting proc file(s)
Collecting procirq file(s)
Collecting dmidecode
Collecting rpm -qa iodrive*
Collecting find /lib/modules
Please send the file /tmp/fio-bugreport-20090921.173256-sdv9ko.tar.bz2
along with your bug report to support@fusionio.com The file is in the /tmp
directory.
```

For example, the filename for a bug report file named

`/tmp/fiobugreport-20090921.173256-sdv9k0.tar.bz2` indicates the following:

- Date (20090921)
- Time (173256, or 17:32:56)
- Misc. information (sdv9ko.tar.bz2)

fio-detach

Description

Detaches and removes the corresponding `/dev/fctx` IO Accelerator block device. The `fio-detach` command waits until the device completes all read/write activity before executing the detach process. The command displays a progress bar and percentage as it completes the process.

NOTE: Before using this utility, be sure that the device you want to detach is not currently mounted and in use.

Syntax

```
fio-detach <device> [options]
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the board number: 0, 1, 2, and so on. For example, `/dev/fct0` indicates the first IO Accelerator installed on the system.

You can specify multiple IO Accelerator devices. For example, `/dev/fct1 /dev/fct2` indicates the second and third IO Accelerator devices installed on the system. You can also use a wildcard to indicate all IO Accelerator devices on the system.

For example, `/dev/fct*`

Options	Description
`-i`	Immediate: Causes a forced immediate detach (does not save metadata). This will fail if the device is in use by the OS.
`-q`	Quiet: Disables the display of the progress bar and percentage.

NOTE: Detaching an IO Accelerator device might fail with an error indicating that the device is busy. This might occur if the IO Accelerator device is part of a software RAID volume (0,1,5), is mounted, or if some process has the device open.
The tools fuser, mount, and lsof can be helpful to determine what is holding the device open.

fio-format

Description

IMPORTANT: The IO Accelerator devices are shipped pre-formatted. `fio-format` is not required except to change the logical size or block size of a device, or to erase user data on a device. To ensure the user data is truly erased, use fio-sure-erase.

The `fio-format` utility performs a low-level format of the board. By default, `fio-format` displays a progress-percentage indicator as it runs.

CAUTION: Use this utility with care since it deletes all user information on the IO Accelerator.

IMPORTANT: Use the `-s` or `-o` option to change the default capacity of the device size. When enabled, the `-s` and `-o` options must include the size or percentage indicators.

NOTE: Use a large block (sector) size to reduce IO Accelerator memory consumption. For example: 4096 bytes. Be aware that some applications are not compatible with non-512-byte sector sizes.

Syntax

```
fio-format [options] <device>
```

where <device> is the name of the device node (/dev/fctx), where x indicates the device number: 0, 1, 2, and so on. For example, /dev/fct0 indicates the first IO Accelerator device installed on the system.

Options	Description
-b <size B|K>	Set the block (sector) size, in bytes or KiBytes (base 2). The default is 512 bytes. For example: -b 512B or -b 4K (B in 512B is optional).
-f	Force the format size, bypassing normal checks and warnings. This option may be needed in rare situations when fio-format does not proceed properly. (The "Are you sure?" prompt still appears unless you use the -y option.) This option can only be used with the -o option.
-q	Quiet mode: Disable the display of the progress-percentage indicator.
-s <size B|K|M|G|T|%>	Set the device capacity as a specific size (in TB, GB, or MB) or as a percentage of the advertised capacity: • T Number of terabytes (TB) to format • G Number of gigabytes (GB) to format • M Number of megabytes (MB) to format • % Percentage, such as 70% (The percent sign must be included)
-o <size B|K|M|G|T|%>	Over-format the device size (to greater than the preset capacity), where the maximum size equals the maximum physical capacity. If a percentage is used, it corresponds to the maximum physical capacity of the device. (Size is required for the -o option; see the -s option above for size indicator descriptions.) Before using this option, for supported recommendations, contact HP Support (http://www.hp.com/support)
-r	Enable fast rescan on non-ordered shutdowns at the cost of some capacity.
-y	Auto-answer "yes" to all queries from the application (bypass prompts).

fio-pci-check

Description

Checks for errors on the PCI bus, specifically for IO Accelerators. This utility displays the current status of each IO Accelerator. It also prints the standard PCI Express error information and resets the state.

It is normal to see a few errors (perhaps as many as five) when fio-pci-check is initially run. Subsequent runs reveal only one or two errors during several hours of operation.

Syntax

```
fio-pci-check [options]
```

Options	Description
-d <value>	1 = Disable the link; 0 = enable the link (Not recommended)
-f	Scan every device in the system
-i	Print the device serial number. This option is invalid when the IO Accelerator is loaded.
-r	Force the link to retrain
-v	Verbose: Print extra data about the hardware

fio-snmp-agentx

Description

This utility is an SNMP subagent that implements the SNMP cpqIODrv-MIB for the IO Accelerator driver. The `fio-snmp-agentx` utility communicates with the SNMP master agent via the agentx protocol.

Syntax

```
fio-snmp-agentx [options]
```

Options	Description
`-f`	Forces the subagent to run in the foreground instead of as a daemon.
`-l`	`<log file>`: Log file to use.
`-s`	Sends errors to `stderr` instead of to `syslog`.

fio-status

Description

Provides detailed information on installed devices. This utility operates on either `fctx` or `fiox` devices. The utility runs from `root` level and having the IO Accelerator loaded. If no IO Accelerator is loaded, there is less content reported from the query.

Enabling `fio-status` provides alerts for certain error modes, such as a minimal-mode, read-only mode, and write-reduced mode, describing what is causing the condition.

Syntax

```
fio-status [<device>] [<options>]
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the card number: `0, 1, 2`, and so on. For example, `/dev/fct0` indicates the first IO Accelerator device installed on the system.

If `<dev>` is not specified, `fio-status` displays information for all cards in the system. If the IO Accelerator is not loaded, this parameter is ignored.

Options	Description
`-a`	Report all available information for each device.
`-e`	Show all errors and warnings for each device. This option is for diagnosing issues, and it hides other information such as format sizes.
`-c`	Count: Report only the number of IO Accelerator devices installed.
`-d`	Show basic information set plus the total amount of data read and written (lifetime data volumes). This option is not necessary when the `-a` option is used.
`-fj`	Format JSON: creates the output in JSON format.
`-fx`	Format XML: creates the output in XML format.
`-u`	Show unavailable fields. Only valid with `-fj` or `-fx`.
`-U`	Show unavailable fields and details why. Only valid with `-fj` or `-fx`.
`-F<field>`	Print the value for a single field (For field names, see the next option). Requires that a device be specified. Multiple `-F` options may be specified.
`-l`	List the fields that can be individually accessed with `-F`.

△ **CAUTION: Output Change**
Starting with version 3.0.0 and later, the standard formatting of `fio-status` output has changed. This will affect any custom management tools that used the output of this utility.

Basic information: If no options are used, `fio-status` reports the following basic information:

- Number and type of devices installed in the system
- IO Accelerator version

Adapter information:

- Adapter type
- Product number
- External power status
- PCIe power limit threshold (if available)
- Connected IO Accelerator devices

Block device information:

- Attach status
- Product name
- Product number
- Serial number
- PCIe address and slot
- Firmware version
- Size of the device, out of total capacity
- Internal temperature (average and maximum, since IO Accelerator load) in degrees Centigrade
- Health status: healthy, nearing wearout, write-reduced or read-only
- Reserve capacity (percentage)
- Warning capacity threshold (percentage)

Data volume information: If the -d option is used, the following data volume information is reported in addition to the basic information:

- Physical bytes written
- Physical bytes read

All information: If the `-a` option is used, all information is printed, which includes the following information in addition to basic and data volume information.

Adapter information:

- Manufacturer number
- Part number
- Date of manufacture
- Power loss protection status
- PCIe bus voltage (average, minimum, maximum)

- PCIe bus current (average, maximum)
- PCIe bus power (average, maximum)
- PCIe power limit threshold (watts)
- PCIe slot available power (watts)
- PCIe negotiated link information (lanes and throughput)

Block device information:

- Manufacturer's code
- Manufacturing date
- Vendor and sub-vendor information
- Format status and sector information (if device is attached)
- FPGA ID and Low-level format GUID
- PCIe slot available power
- PCIe negotiated link information
- Card temperature, in degrees Centigrade
- Internal voltage, average and maximum.
- Auxiliary voltage: average and maximum.
- Percentage of good blocks, data and metadata
- Lifetime data volume statistics
- RAM usage

Error mode information: If the IO Accelerator is in minimal mode, read-only mode, or write-reduced mode when `fio-status` is run, the following differences occur in the output:

- Attach status is Status unknown: Driver is in MINIMAL MODE:
- The reason for the minimal mode state appears (for example, Firmware is out of date. Update firmware.)
- "Geometry and capacity information not available." is displayed.
- No media health information appears.

fio-sure-erase

△ **CAUTION:** Do not use this utility if there are any IO Accelerator devices installed in the system that are not selected to be cleared or purged of data.

- Ensure that you back up any data prior to activation of this utility.
- Remove any devices that are not targeted for purge.
- After the data is removed from the target devices, it is purged.
- There is no recovery to this action.

CAUTION: If the device is in Read-only mode, perform a format using `fio-format` before running `fio-sure-erase`.
The `fio-sure-erase` utility cannot erase the device if it is in Minimal mode. Updating the firmware might move the device out of Minimal Mode. If the device remains in Minimal mode, contact HP Support (http://www.hp.com/support) for assistance.

IMPORTANT: Prior to reactivating the device, format the device with `fio-format` after running `fio-sure-erase`.

To run `fio-sure-erase`, the block device must be detached. For more information, see "`fio-detach` (on page 42)."

Description

The `fio-sure-erase` is a command-line utility that securely removes data from IO Accelerator devices. It complies with the Clear and Purge level of destruction from the following standards:

- DOD 5220.22-M—Complies with instructions for Flash EPROM.
- NIST SP800-88—Complies with instructions for Flash EPROM.

For more information, see the following sections on "Clear support" and "Purge support."

Syntax

```
fio-sure-erase [options] <device>
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the card number: `0, 1, 2,` and so on. For example, `/dev/fct0` indicates the first IO Accelerator device installed on the system. To view this device node, use "`fio-status` (on page 45)."

IMPORTANT: **Products with Multiple Devices**
`fio-sure-erase` applies to individual IO Accelerator devices. For example, if you are planning to purge an ioDrive Duo device, perform this operation on each of the two IO Accelerator devices.

Options	Description
`-p`	Purge instead of Clear: performs a write followed by an erase. For more information, see "Purge support."
`-y`	No confirmation: Does not require a yes/no response to execute the utility.
`-q`	Quiet: Does not display the status bar.

IMPORTANT: If `fio-sure-erase` is run without options, then a Clear is performed. For more information, see "Clear support."

Each block of memory consists of uniform 1 bits or 0 bits.

Clear support

A Clear is the default state of running `fio-sure-erase` (with no options), and refers to the act of performing a full low-level erase (every cell pushed to 1) of the entire NAND media, including retired erase blocks.

Metadata that is required for operation will not be destroyed (media event log, erase counts, physical bytes read/written, performance and thermal history), but any user-specific metadata will be destroyed.

The following describes the steps taken in the Clear operation:

1. Creates a unity map of every addressable block (this allows `fio-sure-erase` to address every block, including previously unmapped bad blocks).
2. For each block, performs an erase cycle (every cell is pushed to 1).
3. Restores the bad block map.

Formats the device (the purpose of this is to make the device usable again, the utility erases all of the headers during the clear).

Purge support

A Purge is implemented by using the `-p` option with `fio-sure-erase`. Purge refers to the act of first overwriting the entire NAND media (including retired erase blocks) with a single character (every cell written to logical 0), and then performing a full chip erase (every cell pushed to 1) across all media (including retired erase blocks).

Metadata that is required for operation will not be destroyed (media event log, erase counts, physical bytes read/written, performance and thermal history), but any user-specific metadata will be destroyed.

The Purge operation includes the following steps:

1. Creates a unity map of every addressable block (this allows `fio-sure-erase` to address every block, including previously unmapped bad blocks).
2. For each block, performs a write cycle (every cell is pushed to 0).
3. For each block, performs an erase cycle (every cell pushed to 1).
4. Restores the bad block map.

Formats the device (the purpose of this is to make the device usable again, the utility erases all of the headers during the clear).

fio-update-iodrive

CAUTION: HP strongly recommends that data is backed up on any IO Accelerator device before performing a firmware upgrade.

Description

Updates the IO Accelerator device's firmware. This utility scans the PCIe bus for all IO Accelerator devices and updates them.

A progress bar and percentage are shown for each device as the update completes.

CAUTION:

- During a firmware upgrade, it is critical to maintain steady power or risk failure of the IO Accelerator device. Connecting a qualified UPS is recommended prior to performing a firmware upgrade.
- It is critical to load the driver after each firmware upgrade step when scheduling sequential, multiple firmware upgrades (example: 1.2.7 to 2.1.0 to 2.3.1). If the driver is not loaded, the on-drive format will not be changed and there will be data loss.
- Data loss may occur if the IO Accelerator device firmware is downgraded. Contact HP Support (http://www.hp.com/support) for recommendations.
- The default to upgrade all IO Accelerator devices does not use the `-d` or `-s` option. The firmware is located in the `<ioaccelerator_version.fff>` file. Confirm that all devices need the firmware upgrade. The `-p` (Pretend) option, can be run to view the possible results of the update.
- Ensure that all IO Accelerator devices are detached before updating the firmware.
- **Upgrade Path**

  There is a specific upgrade path to follow when upgrading an IO Accelerator device. Consult the *Release Notes* for this IO Accelerator release before upgrading any IO Accelerator devices.

IMPORTANT: If you receive an error message when updating the firmware that instructs you to update the midprom information, contact HP Customer Support (http://www.hp.com/support).

To update one or more specific devices:

If the IO Accelerator is loaded, use the `-d` option with the device number.

Syntax

```
fio-update-iodrive [options] <iodrive_version.fff>
```

where `<iodrive_version.fff>` is the path and firmware archive file provided by HP. The default path is `/usr/share/fio/firmware`. This parameter is required.

Options	Description
`-d`	Updates the specified devices (by `fctx`, where `x` is the number of the device shown in `fio-status`). If this option is not specified, all devices are updated. Use the `-d` or `-s` options carefully. Updating the wrong IO Accelerator device could damage that device.
`-f`	Force upgrade (used primarily to downgrade to an earlier firmware version). If the IO Accelerator is not loaded, this option also requires the `-s` option. Use the `-f` option carefully. Updating the wrong IO Accelerator device could damage that device.
`-l`	List the firmware available in the archive.
`-p`	Pretend: Shows what updates would be done. However, the actual firmware is not modified.
`-c`	Clears locks placed on a device.

Options	Description
-q	Runs the update process without displaying the progress bar or percentage.
-y	Confirm all warning messages.
-s	Updates the devices in the specified slots using "*" as a wildcard for devices. The slots are identified in the following PCIe format (as shown in lspci): [[[[<domain>]:]<bus>]:][<slot>][.[<func>]]

All three external LED indicators light up during the update process.

Monitoring IO Accelerator health

NAND flash and component failure

The IO Accelerator is a highly fault-tolerant storage subsystem that provides many levels of protection against component failure and the loss nature of solid state storage. However, as in all storage subsystems, component failures might occur.

By pro-actively monitoring device age and health, you can ensure reliable performance over the intended product life.

Health metrics

The IO Accelerator manages block retirement using pre-determined retirement thresholds. The HP IO Accelerator Management Tool and the `fio-status` utilities show a health indicator that starts at 100 and counts down to 0. As certain thresholds are crossed, various actions are taken.

At the 10% healthy threshold, a one-time warning is issued. For more information, see "Health monitoring techniques."

At 0%, the device is considered unhealthy. It enters write-reduced mode, which somewhat prolongs its lifespan so data can be safely migrated off. In this state the IO Accelerator device behaves normally, except for the reduced write performance.

After the 0% threshold, the device will soon enter read-only mode, and any attempt to write to the IO Accelerator device causes an error. Some filesystems might require special mount options to mount a read-only block device in addition to specifying that the mount must be read-only.

For example, under Linux, `ext3` requires that `-o ro, noload` is used. The `noload` option tells the filesystem to not try and replay the journal.

Consider the read-only mode as a final opportunity to migrate data off the device, as device failure is more likely with continued use.

The IO Accelerator device might enter failure mode. In this case, the device is offline and inaccessible. This can be caused by an internal catastrophic failure, improper firmware upgrade procedures, or device wearout.

The IO Accelerator driver manages LEB retirement via use of pre-determined retirement thresholds. The IO Accelerator Management Tool and the `fio-status` utility show a health indicator that starts at 100 and counts down to 0. As certain thresholds are crossed, various actions are taken.

At the 10% healthy threshold, a one-time warning is issued. For more information, see "Health monitoring techniques."

At 0%, the device is considered unhealthy. It enters write-reduced mode, which somewhat prolongs its lifespan so data can be safely migrated. In this state, the IO Accelerator behaves normally except for the reduced write performance.

At some point after the 0% threshold, the device enters read-only mode. Any attempt to write to the IO Accelerator causes an error. Some file systems might require special mount options to mount a read-only

block device, beyond specifying that the mount must be read-only. For example, under Linux, `ext3` requires that `-o ro, noload` is used. The `noload` option directs the file system not to attempt to replay the journal.

Read-only mode must be considered a final opportunity to migrate data off the device since device failure is more likely with continued use.

The IO Accelerator might enter failure mode. In this case, the device is offline and inaccessible. This can be caused by an internal catastrophic failure, improper firmware upgrade procedures, or device wears out.

IMPORTANT:

- For service or warranty-related questions, contact the company you purchased the device from.
- For products with multiple IO Accelerator devices, these modes are maintained independently for each device.

Health monitoring techniques

fio-status

Output from the `fio-status` utility shows the health percentage and drive state. These items are referenced as `Media status` in the following sample output.

```
Found 3 ioMemory devices in this system
Fusion-io driver version: 3.1.0 build 364
Adapter: Single Adapter
        HP IO Accelerator 1.30TB, Product Number:AJ878B,
SN:1133D0248, FIO SN:1134D9565
...
Media status: Healthy; Reserves: 100.00%, warn at 10.00%; Data: 99.12%
Lifetime data volumes:
   Physical bytes written: 6,423,563,326,064
   Physical bytes read   : 5,509,006,756,312
```

HP IO Accelerator Management Tool: In the Device Report tab, look for the Reserve Space percentage in the right column. The higher the percentage, the healthier the drive is likely to be.

SNMP: On Windows or Linux operating systems, see the corresponding section for details on "Configuring the SNMP master agent."

The following Health Status messages are produced by the `fio-status` utility:

- Healthy
- Read-only
- Reduced-write
- Unknown

About flashback protection technology

Like many other flash devices, NAND flash eventually fails with use. Those failures can be either permanent or temporary. Flashback Protection redundancy is designed to address those IO Accelerator chips that experience permanent failures, and provides additional protection above and beyond ECC for soft failures.

Flashback technology provides a real-time RAID-like redundancy at the chip-level, without sacrificing user capacity or performance for fault tolerance. In general, solutions that use physical RAID schemes for redundancy/protection, must either sacrifice capacity (RAID 1), or performance (RAID 5).

Flashback Protection technology, with self-healing properties, ensures higher performance, minimal failure, and longer endurance than all other flash solutions.

Software RAID and health monitoring

Software RAID stacks are typically designed to detect and mitigate the failure modes of traditional storage media. The IO Accelerator attempts to fail as gracefully as possible, and its new failure mechanisms are compatible with existing software RAID stacks. A drive in write-reduced mode participating in a write-heavy workload is evicted from a RAID group for failure to receive data at a sufficient rate. A drive in read-only mode is evicted when write I/Os are returned from the device as failed. Catastrophic failures are detected and handled just as though they were on traditional storage devices.

Performance and tuning

Introduction to performance and tuning

HP IO Accelerator devices provide high bandwidth and high IOPS and are specifically designed to achieve low latency.

As IO Accelerator devices improve in IOPS and low latency, the device performance may be limited by operating system settings and BIOS configuration. To take advantage of the revolutionary performance of IO Accelerator devices, you might have to tune these settings.

While IO Accelerator devices generally perform well out of the box, this section describes some of the common areas where tuning may help achieve optimal performance.

Disabling DVFS

DVFS is a power management technique that adjusts the CPU voltage and frequency to reduce power consumption by the CPU. These techniques help conserve power and reduce the heat generated by the CPU, but they adversely affect performance while the CPU transitions between low-power and high-performance states.

These power-savings techniques are known to have a negative impact on I/O latency and maximum IOPS. When tuning for maximum performance, you might benefit from reducing or disabling DVSF completely, even though this might increase power consumption.

DVFS, if available, should be configurable as part of your operating systems power management features as well as within your system BIOS interface. Within the operating system and BIOS, DVFS features are often found under the ACPI sections. Consult your computer documentation for details.

Limiting APCI C-states

Newer processors have the ability to go into lower power modes when they are not fully utilized. These idle states are known as ACPI C-states. The C0 state is the normal, full power, operating state. Higher C-states (C1, C2, C3, and so on) are lower power states.

While ACPI C-states save on power, they are known to have a negative impact on I/O latency and maximum IOPS. With each higher C-state, typically more processor functions are limited to save power, and it takes time to restore the processor to the C0 state.

These power savings techniques are known to have a negative impact on I/O latency and maximum IOPS. When tuning for maximum performance you might benefit from limiting the C-states or turning them off completely, even though this might increase power consumption.

If your processor has ACPI C-states available, you can typically limit or disable them in the BIOS interface (sometimes referred to as a Setup Utility). APCI C-states might be part of the ACPI menu. For details, see your computer documentation.

Setting NUMA affinity

Servers with a NUMA (Non-Uniform Memory Access) architecture require special installation instructions in order to maximize ioMemory device performance. These servers include the HP ProLiant DL580 and HP DL980 Servers.

On servers with NUMA architecture, during system boot, the BIOS on some systems will not distribute PCIe slots evenly among the NUMA nodes. Each NUMA node contains multiple CPUs. This imbalanced distribution means that, during high workloads, half or more of the CPUs might remain idle while the rest are 100% utilized. To prevent this imbalance, you must manually assign IO Accelerator devices equally among the available NUMA nodes.

For information on setting NUMA affinity, see "NUMA configuration (on page 57)."

Setting the interrupt handler affinity

Device latency can be affected by placement of interrupts on NUMA systems. HP recommends placing interrupts for a given device on the same NUMA socket that the application is issuing I/O from. If the CPUs on this socket are overwhelmed with user application tasks, in some cases it might benefit performance to move the interrupts to a remote socket to help load balance the system.

Many operating systems will attempt to dynamically place interrupts for you and generally make good decisions. Hand tuning interrupt placement is an advanced option that requires profiling of application performance on any given hardware. For information on how to pin interrupts for a given device to specific CPUs, see your operating system documentation.

NUMA configuration

Introduction to NUMA architecture

Servers with NUMA (Non-Uniform Memory Access) architecture require special installation instructions in order to maximize IO Accelerator device performance. These servers include the HP DL580 and the HP DL980 server.

On servers with NUMA architecture, during system boot, the BIOS on some systems will not distribute PCIe slots evenly among the NUMA nodes. Each NUMA node contains multiple CPUs. This imbalanced distribution means that, during high workloads, half or more of the CPUs will remain idle while the rest are 100% utilized. To prevent this imbalance, you must manually assign IO Accelerator devices equally among the available NUMA nodes.

Configuring the IO Accelerator devices for servers with NUMA architecture requires the use of the `FIO_AFFINTIY` parameter with the `fio-config` utility.

NUMA node override parameter

The `numa_node_override` parameter is a list of `<affinity specification>` couplets that specify the affinity settings of all devices in the system. Each item in the couplet is separated by a colon, and each couplet set is separated by a comma.

Syntax

```
numa_node_override=<affinity specification>[,<affinity specification>...]
```

Where each `<affinity specification>` has the following syntax:

```
<fct-number>:<node-number>
```

Simple example

```
numa_node_override=fct4:1,fct5:0,fct7:2,fct9:3
```

has the effect of creating:

Device	Node/group	Process or affinity
fct4	node 1	all processors in node 1
fct5	node 0	all processors in node 0
fct7	node 2	all processors in node 2
fct9	node 3	all processors in node 3

Advanced configuration example

This example server has four NUMA nodes with eight hyper-threaded cores per node (16 logical processors per node, a total of 64 logical processors in the system). This system also uses the expansion configuration and has 11 PCIe expansion slots. During system boot, the system BIOS assigns PCIe slots 1-6 to NUMA node 2 and PCIe slots 7-11 to NUMA node 0. NUMA nodes 1 and 3 have no assigned PCIe slots. This

configuration creates a load balancing problem in the system when IO Accelerator devices are under heavy traffic. During these periods of high use, half of the processors in the system sit idle while the other half of the processors are 100% utilized, thus limiting the throughput of the IO Accelerator devices.

To avoid this situation, you must manually configure the affinity of the IO Accelerator devices using the `FIO_AFFINITY` configuration parameter to distribute the work load across all NUMA nodes. This parameter overrides the default behavior of the IO Accelerator driver. For more information about the `FIO_AFFINITY` configuration parameter, refer to the syntax explanation below.

Syntax:

The following is an example of how to configure 10 HP IO Accelerator ioDrive Duo devices (each with two IO Accelerator devices) in a HP DL580 G7 system manually as described in the preceding paragraphs. Slot 1 is a Generation 1 PCIe slot, so it is not compatible with an ioDrive Duo device. Therefore you can fill slots 2-11 with ioDrive Duo devices. Because each ioDrive Duo device has two IO Accelerator devices, each ioDrive Duo devices has two device numbers (one for each IO Accelerator device). Each slot has two device numbers.

The following tables list the default BIOS NUMA node assignments.

BIOS assigned NUMA node	PCIe slots	FCT device numbers	Processor Affinity
0	7-11	8,9,13,14,18,19,23,24,28,29	All processors in the node
1	None	None	None
2	2-6	135,136,140,141,145,146,150,151, 155,156	All processors in the node
3	None	None	None

Assigned NUMA node	PCIe slots	FCT device numbers	Processor Affinity
0	7-9	8,9,13,14,18,19	All processors in the node (no hex mask)
1	10-11	23,24,28,29	All processors in the node (no hex mask)
2	2-3	135,136,140,141	All processors in the node (no hex mask)
3	4-6	145,146,150,151,155,156	All processors in the node (no hex mask)

In this example, the BIOS creates a load imbalance by assigning the cards to only two NUMA nodes in the system. To balance the work load, enter the following settings:

Manually configure the VSL driver with these override settings, and then set the `numa_node_override` parameter with the following string:

```
numa_node_override=fct8:0,fct9:0,fct13:0,fct14:0,fct18:0,fct19:0,fct23:1,fct
24:1,fct28:1,fct29:1,fct135:2,fct136:2,fct140:2,fct141:2,fct145:3,fct146:3,f
ct150:3,fct151:3,fct155:3,fct156:3
```

Resources

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website (http://www.hp.com/support).

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

For more information

For additional information, see the following HP websites:

- HP BladeSystem technical resources (http://www.hp.com/go/bladesystem/documentation) (white papers and support documents)
- HP BladeSystem components (http://h18004.www1.hp.com/products/blades/components/c-class-compmatrix.html)
- HP support (http://www.hp.com/support)

NOTE: Before contacting HP customer support, run the IO Accelerator bug reporting tool, and have the report with you when you call. To run the IO Accelerator bug reporting tool, enter the `fio-bugreport` command.

Regulatory information

Safety and regulatory compliance

For safety, environmental, and regulatory information, see *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the HP website (http://www.hp.com/support/Safety-Compliance-EnterpriseProducts).

Turkey RoHS material content declaration

Türkiye Cumhuriyeti: EEE Yönetmeliğine Uygundur

Ukraine RoHS material content declaration

Обладнання відповідає вимогам Технічного регламенту щодо обмеження використання деяких небезпечних речовин в електричному та електронному обладнанні, затвердженого постановою Кабінету Міністрів України від 3 грудня 2008 № 1057

Warranty information

HP ProLiant and X86 Servers and Options (http://www.hp.com/support/ProLiantServers-Warranties)

HP Enterprise Servers (http://www.hp.com/support/EnterpriseServers-Warranties)

HP Storage Products (http://www.hp.com/support/Storage-Warranties)

HP Networking Products (http://www.hp.com/support/Networking-Warranties)

Support and other resources

Before you contact HP

Be sure to have the following information available before you call HP:

- Active Health System log (HP ProLiant Gen8 or later products)

  Download and have available an Active Health System log for 3 days before the failure was detected. For more information, see the *HP iLO 4 User Guide* or *HP Intelligent Provisioning User Guide* on the HP website (http://www.hp.com/go/ilo/docs).
- Onboard Administrator SHOW ALL report (for HP BladeSystem products only)

  For more information on obtaining the Onboard Administrator SHOW ALL report, see the HP website (http://h20000.www2.hp.com/bizsupport/TechSupport/Document.jsp?lang=en&cc=us&objectID=c02843807).
- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Product identification number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

HP contact information

For United States and worldwide contact information, see the Contact HP website (http://www.hp.com/go/assistance).

In the United States:

- To contact HP by phone, call 1-800-334-5144. For continuous quality improvement, calls may be recorded or monitored.
- If you have purchased a Care Pack (service upgrade), see the Support & Drivers website (http://www8.hp.com/us/en/support-drivers.html). If the problem cannot be resolved at the website, call 1-800-633-3600. For more information about Care Packs, see the HP website (http://pro-aq-sama.houston.hp.com/services/cache/10950-0-0-225-121.html).

Customer Self Repair

HP products are designed with many Customer Self Repair (CSR) parts to minimize repair time and allow for greater flexibility in performing defective parts replacement. If during the diagnosis period HP (or HP service

providers or service partners) identifies that the repair can be accomplished by the use of a CSR part, HP will ship that part directly to you for replacement. There are two categories of CSR parts:

- **Mandatory**—Parts for which customer self repair is mandatory. If you request HP to replace these parts, you will be charged for the travel and labor costs of this service.
- **Optional**—Parts for which customer self repair is optional. These parts are also designed for customer self repair. If, however, you require that HP replace them for you, there may or may not be additional charges, depending on the type of warranty service designated for your product.

NOTE: Some HP parts are not designed for customer self repair. In order to satisfy the customer warranty, HP requires that an authorized service provider replace the part. These parts are identified as "No" in the Illustrated Parts Catalog.

Based on availability and where geography permits, CSR parts will be shipped for next business day delivery. Same day or four-hour delivery may be offered at an additional charge where geography permits. If assistance is required, you can call the HP Technical Support Center and a technician will help you over the telephone. HP specifies in the materials shipped with a replacement CSR part whether a defective part must be returned to HP. In cases where it is required to return the defective part to HP, you must ship the defective part back to HP within a defined period of time, normally five (5) business days. The defective part must be returned with the associated documentation in the provided shipping material. Failure to return the defective part may result in HP billing you for the replacement. With a customer self repair, HP will pay all shipping and part return costs and determine the courier/carrier to be used.

For more information about HP's Customer Self Repair program, contact your local service provider. For the North American program, refer to the HP website (http://www.hp.com/go/selfrepair).

Réparation par le client (CSR)

Les produits HP comportent de nombreuses pièces CSR (Customer Self Repair = réparation par le client) afin de minimiser les délais de réparation et faciliter le remplacement des pièces défectueuses. Si pendant la période de diagnostic, HP (ou ses partenaires ou mainteneurs agréés) détermine que la réparation peut être effectuée à l'aide d'une pièce CSR, HP vous l'envoie directement. Il existe deux catégories de pièces CSR:

Obligatoire - Pièces pour lesquelles la réparation par le client est obligatoire. Si vous demandez à HP de remplacer ces pièces, les coûts de déplacement et main d'œuvre du service vous seront facturés.

Facultatif - Pièces pour lesquelles la réparation par le client est facultative. Ces pièces sont également conçues pour permettre au client d'effectuer lui-même la réparation. Toutefois, si vous demandez à HP de remplacer ces pièces, l'intervention peut ou non vous être facturée, selon le type de garantie applicable à votre produit.

REMARQUE: Certaines pièces HP ne sont pas conçues pour permettre au client d'effectuer lui-même la réparation. Pour que la garantie puisse s'appliquer, HP exige que le remplacement de la pièce soit effectué par un Mainteneur Agréé. Ces pièces sont identifiées par la mention "Non" dans le Catalogue illustré.

Les pièces CSR sont livrées le jour ouvré suivant, dans la limite des stocks disponibles et selon votre situation géographique. Si votre situation géographique le permet et que vous demandez une livraison le jour même ou dans les 4 heures, celle-ci vous sera facturée. Pour bénéficier d'une assistance téléphonique, appelez le Centre d'assistance technique HP. Dans les documents envoyés avec la pièce de rechange CSR, HP précise s'il est nécessaire de lui retourner la pièce défectueuse. Si c'est le cas, vous devez le faire dans le délai indiqué, généralement cinq (5) jours ouvrés. La pièce et sa documentation doivent être retournées dans l'emballage fourni. Si vous ne retournez pas la pièce défectueuse, HP se réserve le droit de vous facturer les coûts de remplacement. Dans le cas d'une pièce CSR, HP supporte l'ensemble des frais d'expédition et de retour, et détermine la société de courses ou le transporteur à utiliser.

Pour plus d'informations sur le programme CSR de HP, contactez votre Mainteneur Agrée local. Pour plus d'informations sur ce programme en Amérique du Nord, consultez le site Web HP (http://www.hp.com/go/selfrepair).

Riparazione da parte del cliente

Per abbreviare i tempi di riparazione e garantire una maggiore flessibilità nella sostituzione di parti difettose, i prodotti HP sono realizzati con numerosi componenti che possono essere riparati direttamente dal cliente (CSR, Customer Self Repair). Se in fase di diagnostica HP (o un centro di servizi o di assistenza HP) identifica il guasto come riparabile mediante un ricambio CSR, HP lo spedirà direttamente al cliente per la sostituzione. Vi sono due categorie di parti CSR:

Obbligatorie – Parti che devono essere necessariamente riparate dal cliente. Se il cliente ne affida la riparazione ad HP, deve sostenere le spese di spedizione e di manodopera per il servizio.

Opzionali – Parti la cui riparazione da parte del cliente è facoltativa. Si tratta comunque di componenti progettati per questo scopo. Se tuttavia il cliente ne richiede la sostituzione ad HP, potrebbe dover sostenere spese addizionali a seconda del tipo di garanzia previsto per il prodotto.

NOTA: alcuni componenti HP non sono progettati per la riparazione da parte del cliente. Per rispettare la garanzia, HP richiede che queste parti siano sostituite da un centro di assistenza autorizzato. Tali parti sono identificate da un "No" nel Catalogo illustrato dei componenti.

In base alla disponibilità e alla località geografica, le parti CSR vengono spedite con consegna entro il giorno lavorativo seguente. La consegna nel giorno stesso o entro quattro ore è offerta con un supplemento di costo solo in alcune zone. In caso di necessità si può richiedere l'assistenza telefonica di un addetto del centro di supporto tecnico HP. Nel materiale fornito con una parte di ricambio CSR, HP specifica se il cliente deve restituire dei componenti. Qualora sia richiesta la resa ad HP del componente difettoso, lo si deve spedire ad HP entro un determinato periodo di tempo, generalmente cinque (5) giorni lavorativi. Il componente difettoso deve essere restituito con la documentazione associata nell'imballo di spedizione fornito. La mancata restituzione del componente può comportare la fatturazione del ricambio da parte di HP. Nel caso di riparazione da parte del cliente, HP sostiene tutte le spese di spedizione e resa e sceglie il corriere/vettore da utilizzare.

Per ulteriori informazioni sul programma CSR di HP contattare il centro di assistenza di zona. Per il programma in Nord America fare riferimento al sito Web HP (http://www.hp.com/go/selfrepair).

Customer Self Repair

HP Produkte enthalten viele CSR-Teile (Customer Self Repair), um Reparaturzeiten zu minimieren und höhere Flexibilität beim Austausch defekter Bauteile zu ermöglichen. Wenn HP (oder ein HP Servicepartner) bei der Diagnose feststellt, dass das Produkt mithilfe eines CSR-Teils repariert werden kann, sendet Ihnen HP dieses Bauteil zum Austausch direkt zu. CSR-Teile werden in zwei Kategorien unterteilt:

Zwingend – Teile, für die das Customer Self Repair-Verfahren zwingend vorgegeben ist. Wenn Sie den Austausch dieser Teile von HP vornehmen lassen, werden Ihnen die Anfahrt- und Arbeitskosten für diesen Service berechnet.

Optional – Teile, für die das Customer Self Repair-Verfahren optional ist. Diese Teile sind auch für Customer Self Repair ausgelegt. Wenn Sie jedoch den Austausch dieser Teile von HP vornehmen lassen möchten, können bei diesem Service je nach den für Ihr Produkt vorgesehenen Garantiebedingungen zusätzliche Kosten anfallen.

HINWEIS: Einige Teile sind nicht für Customer Self Repair ausgelegt. Um den Garantieanspruch des Kunden zu erfüllen, muss das Teil von einem HP Servicepartner ersetzt werden. Im illustrierten Teilekatalog sind diese Teile mit „No“ bzw. „Nein“ gekennzeichnet.

CSR-Teile werden abhängig von der Verfügbarkeit und vom Lieferziel am folgenden Geschäftstag geliefert. Für bestimmte Standorte ist eine Lieferung am selben Tag oder innerhalb von vier Stunden gegen einen Aufpreis verfügbar. Wenn Sie Hilfe benötigen, können Sie das HP technische Support Center anrufen und sich von einem Mitarbeiter per Telefon helfen lassen. Den Materialien, die mit einem CSR-Ersatzteil geliefert werden, können Sie entnehmen, ob das defekte Teil an HP zurückgeschickt werden muss. Wenn es erforderlich ist, das defekte Teil an HP zurückzuschicken, müssen Sie dies innerhalb eines vorgegebenen Zeitraums tun, in der Regel innerhalb von fünf (5) Geschäftstagen. Das defekte Teil muss mit der zugehörigen Dokumentation in der Verpackung zurückgeschickt werden, die im Lieferumfang enthalten ist. Wenn Sie das defekte Teil nicht zurückschicken, kann HP Ihnen das Ersatzteil in Rechnung stellen. Im Falle von Customer Self Repair kommt HP für alle Kosten für die Lieferung und Rücksendung auf und bestimmt den Kurier-/Frachtdienst.

Weitere Informationen über das HP Customer Self Repair Programm erhalten Sie von Ihrem Servicepartner vor Ort. Informationen über das CSR-Programm in Nordamerika finden Sie auf der HP Website unter (http://www.hp.com/go/selfrepair).

Reparaciones del propio cliente

Los productos de HP incluyen muchos componentes que el propio usuario puede reemplazar (*Customer Self Repair*, CSR) para minimizar el tiempo de reparación y ofrecer una mayor flexibilidad a la hora de realizar sustituciones de componentes defectuosos. Si, durante la fase de diagnóstico, HP (o los proveedores o socios de servicio de HP) identifica que una reparación puede llevarse a cabo mediante el uso de un componente CSR, HP le enviará dicho componente directamente para que realice su sustitución. Los componentes CSR se clasifican en dos categorías:

- **Obligatorio:** componentes para los que la reparación por parte del usuario es obligatoria. Si solicita a HP que realice la sustitución de estos componentes, tendrá que hacerse cargo de los gastos de desplazamiento y de mano de obra de dicho servicio.
- **Opcional:** componentes para los que la reparación por parte del usuario es opcional. Estos componentes también están diseñados para que puedan ser reparados por el usuario. Sin embargo, si precisa que HP realice su sustitución, puede o no conllevar costes adicionales, dependiendo del tipo de servicio de garantía correspondiente al producto.

NOTA: Algunos componentes no están diseñados para que puedan ser reparados por el usuario. Para que el usuario haga valer su garantía, HP pone como condición que un proveedor de servicios autorizado realice la sustitución de estos componentes. Dichos componentes se identifican con la palabra "No" en el catálogo ilustrado de componentes.

Según la disponibilidad y la situación geográfica, los componentes CSR se enviarán para que lleguen a su destino al siguiente día laborable. Si la situación geográfica lo permite, se puede solicitar la entrega en el mismo día o en cuatro horas con un coste adicional. Si precisa asistencia técnica, puede llamar al Centro de asistencia técnica de HP y recibirá ayuda telefónica por parte de un técnico. Con el envío de materiales para la sustitución de componentes CSR, HP especificará si los componentes defectuosos deberán devolverse a HP. En aquellos casos en los que sea necesario devolver algún componente a HP, deberá hacerlo en el periodo de tiempo especificado, normalmente cinco días laborables. Los componentes defectuosos deberán devolverse con toda la documentación relacionada y con el embalaje de envío. Si no enviara el componente defectuoso requerido, HP podrá cobrarle por el de sustitución. En el caso de todas

sustituciones que lleve a cabo el cliente, HP se hará cargo de todos los gastos de envío y devolución de componentes y escogerá la empresa de transporte que se utilice para dicho servicio.

Para obtener más información acerca del programa de Reparaciones del propio cliente de HP, póngase en contacto con su proveedor de servicios local. Si está interesado en el programa para Norteamérica, visite la página web de HP siguiente (http://www.hp.com/go/selfrepair).

Customer Self Repair

Veel onderdelen in HP producten zijn door de klant zelf te repareren, waardoor de reparatieduur tot een minimum beperkt kan blijven en de flexibiliteit in het vervangen van defecte onderdelen groter is. Deze onderdelen worden CSR-onderdelen (Customer Self Repair) genoemd. Als HP (of een HP Service Partner) bij de diagnose vaststelt dat de reparatie kan worden uitgevoerd met een CSR-onderdeel, verzendt HP dat onderdeel rechtstreeks naar u, zodat u het defecte onderdeel daarmee kunt vervangen. Er zijn twee categorieën CSR-onderdelen:

Verplicht: Onderdelen waarvoor reparatie door de klant verplicht is. Als u HP verzoekt deze onderdelen voor u te vervangen, worden u voor deze service reiskosten en arbeidsloon in rekening gebracht.

Optioneel: Onderdelen waarvoor reparatie door de klant optioneel is. Ook deze onderdelen zijn ontworpen voor reparatie door de klant. Als u echter HP verzoekt deze onderdelen voor u te vervangen, kunnen daarvoor extra kosten in rekening worden gebracht, afhankelijk van het type garantieservice voor het product.

OPMERKING: Sommige HP onderdelen zijn niet ontwikkeld voor reparatie door de klant. In verband met de garantievoorwaarden moet het onderdeel door een geautoriseerde Service Partner worden vervangen. Deze onderdelen worden in de geïllustreerde onderdelencatalogus aangemerkt met "Nee".

Afhankelijk van de leverbaarheid en de locatie worden CSR-onderdelen verzonden voor levering op de eerstvolgende werkdag. Levering op dezelfde dag of binnen vier uur kan tegen meerkosten worden aangeboden, indien dit mogelijk is gezien de locatie. Indien assistentie gewenst is, belt u een HP Service Partner om via de telefoon technische ondersteuning te ontvangen. HP vermeldt in de documentatie bij het vervangende CSR-onderdeel of het defecte onderdeel aan HP moet worden geretourneerd. Als het defecte onderdeel aan HP moet worden teruggezonden, moet u het defecte onderdeel binnen een bepaalde periode, gewoonlijk vijf (5) werkdagen, retourneren aan HP. Het defecte onderdeel moet met de bijbehorende documentatie worden geretourneerd in het meegeleverde verpakkingsmateriaal. Als u het defecte onderdeel niet terugzendt, kan HP u voor het vervangende onderdeel kosten in rekening brengen. Bij reparatie door de klant betaalt HP alle verzendkosten voor het vervangende en geretourneerde onderdeel en kiest HP zelf welke koerier/transportonderneming hiervoor wordt gebruikt.

Neem contact op met een Service Partner voor meer informatie over het Customer Self Repair programma van HP. Informatie over Service Partners vindt u op de HP website (http://www.hp.com/go/selfrepair).

Reparo feito pelo cliente

Os produtos da HP são projetados com muitas peças para reparo feito pelo cliente (CSR) de modo a minimizar o tempo de reparo e permitir maior flexibilidade na substituição de peças com defeito. Se, durante o período de diagnóstico, a HP (ou fornecedores/parceiros de serviço da HP) concluir que o reparo pode ser efetuado pelo uso de uma peça CSR, a peça de reposição será enviada diretamente ao cliente. Existem duas categorias de peças CSR:

Obrigatória – Peças cujo reparo feito pelo cliente é obrigatório. Se desejar que a HP substitua essas peças, serão cobradas as despesas de transporte e mão-de-obra do serviço.

Opcional – Peças cujo reparo feito pelo cliente é opcional. Essas peças também são projetadas para o reparo feito pelo cliente. No entanto, se desejar que a HP as substitua, pode haver ou não a cobrança de taxa adicional, dependendo do tipo de serviço de garantia destinado ao produto.

OBSERVAÇÃO: Algumas peças da HP não são projetadas para o reparo feito pelo cliente. A fim de cumprir a garantia do cliente, a HP exige que um técnico autorizado substitua a peça. Essas peças estão identificadas com a marca "No" (Não), no catálogo de peças ilustrado.

Conforme a disponibilidade e o local geográfico, as peças CSR serão enviadas no primeiro dia útil após o pedido. Onde as condições geográficas permitirem, a entrega no mesmo dia ou em quatro horas pode ser feita mediante uma taxa adicional. Se precisar de auxílio, entre em contato com o Centro de suporte técnico da HP para que um técnico o ajude por telefone. A HP especifica nos materiais fornecidos com a peça CSR de reposição se a peça com defeito deve ser devolvida à HP. Nos casos em que isso for necessário, é preciso enviar a peça com defeito à HP dentro do período determinado, normalmente cinco (5) dias úteis. A peça com defeito deve ser enviada com a documentação correspondente no material de transporte fornecido. Caso não o faça, a HP poderá cobrar a reposição. Para as peças de reparo feito pelo cliente, a HP paga todas as despesas de transporte e de devolução da peça e determina a transportadora/serviço postal a ser utilizado.

Para obter mais informações sobre o programa de reparo feito pelo cliente da HP, entre em contato com o fornecedor de serviços local. Para o programa norte-americano, visite o site da HP (http://www.hp.com/go/selfrepair).

カスタマーセルフリペア

修理時間を短縮し、故障部品の交換における高い柔軟性を確保するために、HP製品には多数のCSR部品があります。診断の際に、CSR部品を使用すれば修理ができるとHP（HPまたはHP正規保守代理店）が判断した場合、HPはその部品を直接、お客様に発送し、お客様に交換していただきます。CSR部品には以下の2通りがあります。

- **必須 -** カスタマーセルフリペアが必須の部品。当該部品について、もしもお客様がHPに交換作業を依頼される場合には、その修理サービスに関する交通費および人件費がお客様に請求されます。
- **任意 -** カスタマーセルフリペアが任意である部品。この部品もカスタマーセルフリペア用です。当該部品について、もしもお客様がHPに交換作業を依頼される場合には、お買い上げの製品に適用される保証サービス内容の範囲内においては、別途費用を負担していただくことなく保証サービスを受けることができます。

注： HP製品の一部の部品は、カスタマーセルフリペア用ではありません。製品の保証を継続するためには、HPまたはHP正規保守代理店による交換作業が必須となります。部品カタログには、当該部品がカスタマーセルフリペア除外品である旨が記載されています。

部品供給が可能な場合、地域によっては、CSR部品を翌営業日に届くように発送します。また、地域によっては、追加費用を負担いただくことにより同日または4時間以内に届くように発送することも可能な場合があります。サポートが必要なときは、HPの修理受付窓口に電話していただければ、技術者が電話でアドバイスします。交換用のCSR部品または同梱物には、故障部品をHPに返送する必要があるかどうかが表示されています。故障部品をHPに返送する必要がある場合は、指定期限内（通常は5営業日以内）に故障部品をHPに返送してください。故障部品を返送する場合は、届いた時の梱包箱に関連書類とともに入れてください。故障部品を返送しない場合、HPから部品費用が請求されます。カスタマーセルフリペアの際には、HPは送料および部品返送費を全額負担し、使用する宅配便会社や運送会社を指定します。

客户自行维修

HP 产品提供许多客户自行维修 (CSR) 部件，以尽可能缩短维修时间和在更换缺陷部件方面提供更大的灵活性。如果在诊断期间 HP（或 HP 服务提供商或服务合作伙伴）确定可以通过使用 CSR 部件完成维修，HP 将直接把该部件发送给您进行更换。有两类 CSR 部件：

- **强制性的** — 要求客户必须自行维修的部件。如果您请求 HP 更换这些部件，则必须为该服务支付差旅费和人工费用。
- **可选的** — 客户可以选择是否自行维修的部件。这些部件也是为客户自行维修设计的。不过，如果您要求 HP 为您更换这些部件，则根据为您的产品指定的保修服务类型，HP 可能收取或不再收取任何附加费用。

注：某些 HP 部件的设计并未考虑客户自行维修。为了满足客户保修的需要，HP 要求授权服务提供商更换相关部件。这些部件在部件图解目录中标记为“否”。

CSR 部件将在下一个工作日发运（取决于备货情况和允许的地理范围）。在允许的地理范围内，可在当天或四小时内发运，但要收取额外费用。如果需要帮助，您可以致电 HP 技术支持中心，将会有技术人员通过电话为您提供帮助。HP 会在随更换的 CSR 部件发运的材料中指明是否必须将有缺陷的部件返还给 HP。如果要求您将有缺陷的部件返还给 HP，那么您必须在规定期限内（通常是五 (5) 个工作日）将缺陷部件发给 HP。有缺陷的部件必须随所提供的发运材料中的相关文件一起返还。如果未能送还有缺陷的部件，HP 可能会要求您支付更换费用。客户自行维修时，HP 将承担所有相关运输和部件返回费用，并指定快递商/承运商。

有关 HP 客户自行维修计划的详细信息，请与您当地的服务提供商联系。有关北美地区的计划，请访问 HP 网站 (http://www.hp.com/go/selfrepair)。

客戶自行維修

HP 產品設計了許多「客戶自行維修」(CSR) 的零件以減少維修時間，並且使得更換瑕疵零件時能有更大的彈性。如果在診斷期間 HP（或 HP 服務供應商或維修夥伴）辨認出此項維修工作可以藉由使用 CSR 零件來完成，則 HP 將直接寄送該零件給您作更換。CSR 零件分為兩種類別：

- **強制的** — 客戶自行維修所使用的零件是強制性的。如果您要求 HP 更換這些零件，HP 將會向您收取此服務所需的外出費用與勞動成本。
- **選購的** — 客戶自行維修所使用的零件是選購的。這些零件也設計用於客戶自行維修之用。不過，如果您要求 HP 為您更換，則可能需要也可能不需要負擔額外的費用，端視針對此產品指定的保固服務類型而定。

備註：某些 HP 零件沒有消費者可自行維修的設計。為符合客戶保固，HP 需要授權的服務供應商更換零件。這些零件在圖示的零件目錄中，被標示為「否」。

基於材料取得及環境允許的情況下，CSR 零件將於下一個工作日以快遞寄送。在環境的允許下當天或四小時內送達，則可能需要額外的費用。若您需要協助，可致電「HP 技術支援中心」，會有一位技術人員透過電話來協助您。不論損壞的零件是否必須退回，HP 皆會在與 CSR 替換零件一起運送的材料中註明。若要將損壞的零件退回 HP，您必須在指定的一段時間內（通常為五 (5) 個工作天），將損壞的零件寄回 HP。損壞的零件必須與寄送資料中隨附的相關技術文件一併退還。如果無法退還損壞的零件，HP 可能要向您收取替換費用。針對客戶自行維修情形，HP 將負責所有運費及零件退還費用並指定使用何家快遞/貨運公司。

如需 HP 的「客戶自行維修」方案詳細資訊，請連絡您當地的服務供應商。至於北美方案，請參閱 HP 網站 (http://www.hp.com/go/selfrepair)。

고객 셀프 수리

HP 제품은 수리 시간을 최소화하고 결함이 있는 부품 교체 시 더욱 융통성을 발휘할 수 있도록 하기 위해 고객 셀프 수리(CSR) 부품을 다량 사용하여 설계되었습니다. 진단 기간 동안 HP(또는 HP 서비스 공급업체 또는 서비스 협력업체)에서 CSR 부품을 사용하여 수리가 가능하다고 판단되면 HP는 해당 부품을 바로 사용자에게 보내어 사용자가 교체할 수 있도록 합니다. CSR 부품에는 두 가지 종류가 있습니다.

- **고객 셀프 수리가 의무 사항인 필수 부품.** 사용자가 HP에 이 부품의 교체를 요청할 경우 이 서비스에 대한 출장비 및 작업비가 청구됩니다.
- **고객 셀프 수리가 선택 사항인 부품.** 이 부품들도 고객 셀프 수리가 가능하도록 설계되었습니다. 하지만 사용자가 HP에 이 부품의 교체를 요청할 경우 사용자가 구입한 제품에 해당하는 보증 서비스 유형에 따라 추가 비용 없이 교체가 가능할 수 있습니다.

참고: 일부 HP 부품은 고객 셀프 수리가 불가능하도록 설계되었습니다. HP는 만족스러운 고객 보증을 위해 공인 서비스 제공업체를 통해 부품을 교체하도록 하고 있습니다. 이러한 부품들은 Illustrated Parts Catalog에 "No"라고 표시되어 있습니다.

CSR 부품은 재고 상태와 지리적 조건이 허용하는 경우 다음 영업일 납품이 가능하도록 배송이 이루어집니다. 지리적 조건이 허용하는 경우 추가 비용이 청구되는 조건으로 당일 또는 4시간 배송이 가능할 수도 있습니다. 도움이 필요하시면 HP 기술 지원 센터로 전화하십시오. 전문 기술자가 전화로 도움을 줄 것입니다. HP는 결함이 발생한 부품을 HP로 반환해야 하는지 여부를 CSR 교체 부품과 함께 배송된 자료에 지정합니다. 결함이 발생한 부품을 HP로 반환해야 하는 경우에는 지정된 기간 내(통상 영업일 기준 5일)에 HP로 반환해야 합니다. 이 때 결함이 발생한 부품은 제공된 포장 재료에 넣어 관련 설명서와 함께 반환해야 합니다. 결함이 발생한 부품을 반환하지 않는 경우 HP가 교체 부품에 대해 비용을 청구할 수 있습니다. 고객 셀프 수리의 경우, HP는 모든 운송 및 부품 반환 비용을 부담하며 이용할 운송업체 및 택배 서비스를 결정합니다.

HP 고객 셀프 수리 프로그램에 대한 자세한 내용은 가까운 서비스 제공업체에 문의하십시오. 북미 지역의 프로그램에 대해서는 HP 웹 사이트(http://www.hp.com/go/selfrepair)를 참조하십시오.

Acronyms and abbreviations

ACPI

Advanced Configuration and Power Interface

DVFS

dynamic voltage and frequency scaling

IOPS

input/output operations per second

LEB

Logical Erase Block

LVM

Logical Volume Manager

MIB

management information base

NAND

Not AND

NUMA

Non-Uniform Memory Architecture

PSP

HP ProLiant Support Pack

RFC

request for comments

RHEL

Red Hat Enterprise Linux

RPM

Red Hat Package Manager

SMH

System Management Homepage

VSL

virtual storage layer

Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (mailto:docsfeedback@hp.com). Include the document title and part number, version number, or the URL when submitting your feedback.

Index

A

about this guide 6
advanced NUMA configuration example 57
authorized reseller 61

B

battery replacement notice 60
building a RAID 10 configuration across multiple IO Accelerator Duos 25
building driver from source package 14
building RPM package 14

C

command-line utilities 35
common maintenance tasks 37
configuring RAID 23
configuring the SNMP master agent 28
configuring the SNMP subagent 29
contents summary 6
controlling driver loading 18
CSR (customer self repair) 61
customer self repair (CSR) 61

D

disabling auto attach 38
disabling DVFS 55
disabling the driver 39
documentation feedback 71
driver unloads 20

E

Enabling PCIe power 22
enabling PCIe power override 36
enabling SNMP test mode 30
enabling the override parameter 37
European Union notice 60

F

filesystems, mounting 20
fio-attach utility 40
fio-beacon utility 41
fio-bugreport utility 41
fio-detach utility 42
fio-format utility 43
fio-pci-check utility 44
fio-snmp-agentx utility 45
fio-status utility 45
fio-sure-erase utility 47
fio-update-iodrive utility 49
firmware, upgrading 15, 21
flashback protection technology 54

H

health metrics 52
health monitoring techniques 53
HP IO Accelerator Management Tool 35
HP Subscriber's Choice for Business 59
HP technical support 61
HP, contacting 61

I

installing the SNMP subagent 28
installing, RPM package 12
introduction 7

L

launching the SNMP master agent 27
LED indicators 35
limiting APCI C-states 55
loading the driver 18
Logical Volume Manager 22

M

maintenance 20, 35, 37, 38, 39
maintenance tools 35
making an array persistent 24
manually running the SNMP subagent 29
module parameters 20
monitoring IO Accelerator health 52
mounting filesystems 20

N

NAND flash and component failure 52
NUMA architecture 57
NUMA configuration 57
NUMA node override parameter 57

O

one-time configuration 21
operating environment 9
overview 7

P

PCIe power override, enabling 36
PCIe power, enabling 22
performance and tuning 55
performance and tuning, introduction 55
persistent configuration 21
phone numbers 61

R

RAID 0/Striped 23
RAID 1/Mirrored 25
RAID 10 25
RAID configuration 23
regulatory compliance identification numbers 60
regulatory compliance notices 60
required information 61
resources 59, 61
RPM package 12
running the SNMP subagent 29

S

series number 60
setting driver options 20
setting interrupt handler affinity 56
setting NUMA affinity 56
SNMP agentx subagent 28
SNMP files and directories 27
SNMP master agent 27
SNMP MIB support 33
SNMP sample config files, 30
SNMP, Linux details 27
software installation 12
software RAID and health monitoring 54
subagent log file 29
support 61
supported firmware revisions 9
supported hardware 9

T

Taiwan battery recycling notice 60
technical support 61
telephone numbers 61
Trim support 26
troubleshooting SNMP 33

U

uninstalling driver RPM package 38
uninstalling utilities 38
unloading the driver 38
unmanaged shutdown issues 39
upgrading firmware 15, 21
upgrading, procedure 16
using Discard 26
using init scripts 19
using the IO Accelerator as a swap 22
using the Logical Volume Manager 22
utilities 40
utilities reference 40